



(51) International Patent Classification:

H04W 4/00 (2009.01) *H04L* 12/24 (2006.01)
G06F 17/30 (2006.01) *H04L* 29/06 (2006.01)
H04L 12/28 (2006.01)

(21) International Application Number:

PCT/US2017/029341

(22) International Filing Date:

25 April 2017 (25.04.2017)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/326,894 25 April 2016 (25.04.2016) US

(71) Applicant: **CONVIDA WIRELESS, LLC** [US/US]; 200
Bellevue Parkway, Suite 300, Wilmington, DE 19809-3727
(US).

(72) Inventors: **SINGH, Sneha, Aman**; 7977 170th Avenue
NE, Apt. B620, Redmond, WA 98052 (US). **LI, Xu**; 200
Bellevue Parkway, Suite 300, Wilmington, DE 19809-3727
(US). **LY, Quang**; 200 Bellevue Parkway, Suite 300, Wilm-
ington, DE 19809-3727 (US). **STERNBERG, Gregory, S.**;
43 Sorrel Run, Mt. Laurel, NJ 08054 (US). **RAHMAN,**
Shamim, Akbar; 200 Bellevue Parkway, Suite 300, Wilm-
ington, DE 19809-3727 (US). **WANG, Chonggang**; 200
Bellevue Parkway, Suite 300, Wilmington, DE 19809-3727
(US).

(74) Agent: **O'MALLEY, Joseph P.** et al.; Cira Centre, 12th
Floor, 2929 Arch Street, Philadelphia, PA 19104-2891
(US).

(81) Designated States (unless otherwise indicated, for every
kind of national protection available): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ,
CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR,

(54) Title: DATA STREAM ANALYTICS AT SERVICE LAYER




FIG. 22

(57) **Abstract:** A modular and distributed architecture for data stream processing and analysis is described to incorporate data stream analytics capabilities, called Data Stream Analytics Service (DSAS) in the IoT/M2M service layer. Each service layer node hosting DSAS can be split into two independent modules, Stream Forwarder and Stream Analytics Engine. Stream Forwarder is a light weight processing modules that can be responsible for data preprocessing and routing. Stream Analytics Engine is responsible for performing actual analytics on the data stream. Separating the two functionalities enables the service layer nodes to efficiently distribute stream analytics tasks across multiple nodes.



KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

DATA STREAM ANALYTICS AT SERVICE LAYER

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This Application claims the benefit of U.S. Provisional Patent Application Serial No. 62/326,894, filed April 25, 2016, the disclosure of which is hereby incorporated by reference as if set forth in its entirety.

BACKGROUND

[0002] A data stream is a continuous and dynamic sequence of data (information) in the process of being transmitted between nodes across interconnected communication channels in the form of events, messages, and so on. A data stream is typically data that has not yet been stored in storage such as disk memory. A data stream can be massive and is potentially unbounded in size, time-varying, and in many cases, of high throughput.

[0003] A good example of massive, high throughput and dynamic data is Internet traffic data. As per the Internet Live Stats website, a few thousands of petabytes of Internet traffic data is generated every day. Some other examples of data streams are telecommunication data, satellite data, stock market/financial trade data.

[0004] Internet of Things/ Machine to Machine (IoT/M2M) applications generate a large amount of data streams as well. With the increasing number of “things / machines” in IoT/M2M systems, the volume of the streaming data to and from the IoT/M2M devices is increasing exponentially. Figure 1 shows connected devices in a smart home setup. These devices continually generate data and communicate with other devices in the system by transmitting/receiving data, leading to continuous generation of large amounts of data streams from distributed sources and devices.

[0005] A data stream ‘S’ can be denoted as a sequence of tuples, where each tuple is of the format $\langle a_1, a_2, \dots, a_n \rangle$ where ‘ a_i ’ is the i-th attribute (or property) of the tuple. Figure 2 shows an example data stream generated by a Global Positioning System (GPS) probe device such as a GPS tracking vehicle. Each line in the figure represents one tuple, where each tuple comprises of six attributes: timestamp, validity of GPS data (‘A’ denotes that the data is valid), Latitude, Longitude, Altitude, and Speed. (A GPS data stream tuple generally contains many more attributes than shown in Figure 2. The GPS data stream tuple illustrated in Figure 2 has been shortened for simplicity.)

[0006] In many applications, making sense of data while it is still in the process of streaming is of great importance so that necessary actions or measures can be taken in real time based on the information retrieved from the data stream.

[0007] With the exponential growth of streaming data, it is becoming increasingly important to derive meaningful insights from the data in real time. Data stream analytics (or stream analytics or streaming data analytics) is the process of extracting and deriving useful information from continuous and streaming data in real time. Stream analytics include, but is not limited to, detecting or identifying events, patterns or trends from the data streams, computing statistical aggregation over the data streams, and performing correlation and predictive analysis over the data streams. One example application where real time stream analytics finds great usage is the domain of network security where identifying and predicting network attacks in real time can avoid a node or a network of nodes from being compromised. Some other important example applications of data stream analytics are in stock price prediction, dynamic network performance monitoring and tuning, and real time fraud detection.

[0008] With the evolution of IoT/M2M and the growing dependence of the world on sensors and devices, there is also a growing need to analyze the rich data generated by these devices, almost as fast as they are generated. For most of this data, the real value lies in exploiting the information hidden in the data in real time and taking instantaneous actions based on the insights provided by the data. Hence, real time data analytics is starting to become an integral part of IoT/M2M systems.

[0009] Consider the example of an Intelligent Intersection Traffic System (IITS) which is a framework for efficient and safe management of traffic at the intersection. An efficient implementation of IITS can reduce the traffic collisions at an intersection drastically by taking real time decisions depending on the traffic conditions at the intersection, and generating necessary real time alerts.

[0010] Due to the dynamic nature of the data stream, i.e. the time-varying properties of the data stream, processing and analyzing the stream using traditional methods poses several challenges. Analytics is challenging over dynamic data for a number of reasons.

[0011] On-the-fly processing – The most fundamental concept of the data stream is that it is in a continuous streaming state, and not stored in any disk or database during the time of its observation. This makes multiple passes over a data stream, i.e. backtracking a data stream, almost impossible. Hence, all the necessary information required for streaming analytics need to be retrieved from the data in a single pass in an incremental fashion.

[0012] Massive size of data stream, small processing time and limited main memory constraint – Due to the large, possibly unbounded, size of the data stream, it is generally practically infeasible to temporarily store the entire data stream in disk for real time analysis. In order to perform real time analysis, the data stream needs to be processed and analyzed within a small processing time (depending on the application). Further, the additional requirement of having to analyze the data in real time using a very small processing time, makes it even more infeasible to store the data in the disk for analysis. The overhead of reading data from and/or writing data into the disk (also called disk Input/Output or disk I/O) is very high and adds significant delay to the processing time of data when compared to processing data in main memory (for instance, in Random Access Memory or RAM), also called in-memory processing of data. This additional overhead renders the storage of data unsuitable for real time analysis.

[0013] Further, due to the massive size of the data stream and limited size of the main memory, it is practically not feasible to store the entire data stream in main memory. The size of main memory today is generally in the order of Gigabytes, or GB, (10^9 bytes) while the size of the data stream can be in the order of Petabytes, or PB, (10^6 GB), or more. Hence, in many cases, only the important information from the data, required for the purpose of analytics, is extracted and stored in main memory in a concise fashion.

[0014] High throughput of data stream – In many applications, the rate of generation of data can be very high, on the order of Terabytes/second (TB/s) or more. Processing data of such high throughput in a small window of time adds additional challenge to stream analytics.

[0015] Due to the dynamic and transient nature of the data stream, it is not practically feasible to use a traditional Database Management System (DBMS) or other traditional batch processing system for stream processing/analysis. The main challenge of using a traditional DBMS for real time analytics arises from the fact that it is not practically feasible to locally store a data stream on a disk or backtrack over a data stream. Hence, the data stream needs to be processed and analyzed in a single pass, unlike in DBMS where stored data can be backtracked over as many times as possible. The single pass constraint over a data stream also implies queries pertaining to the data analysis needs to be long running, continuous and persistent so that the data stream is incrementally analyzed as soon as it is observed. This is contrary to the traditional DBMS where data is generally static and hence the queries can be ad hoc. Table 1 shows the key differences in requirements for processing data stream as compared to the disk resident data (stored in disk).

Table 1: Difference in Processing Disk Resident Data *vs* Data Stream

Disk Resident Data	Data Stream
Stored on disk	Not stored to disk
Support for ad hoc one time queries	Support for persistent and continuous set of pre-defined queries
Random / sequential access of data, unlimited reads	Sequential access of data in a single pass
Processed using large disk space	Processed using small and limited memory space

[0016] Figure 3 shows a general logical framework of Data Stream Processing System (DSPS) 302 or Data Stream Management System (DSMS). This system has a processing engine 304 that continuously manages and processes a data stream. The infeasibility of storing the entire data stream in memory suggests summarizing the observed data stream and storing only important information from the stream in a concise way. We refer to the concise information extracted from the data to summarize the data stream as stream synopsis or stream state. A stream state is generally much smaller in size as compared to the size of the data stream. Hence, the stream state can be stored in memory for faster access.

[0017] The DSPS 302 generates stream state depending on the queries running in the DSPS 302. These queries, as discussed above, are long standing and continuous, i.e. the queries need to be predefined so that DSPS generates the corresponding synopsis accordingly. Since the stream state is only the summary of an entire stream, the answers to the query may be approximate. Generally in the case of approximate queries, users or applications interested in deriving results from such queries are given a resource vs. accuracy tradeoff. The larger the size of the memory (and/or processing capacity) allocated for the analytical operation pertaining to the query, the higher the answer.

[0018] In order to understand the data stream processing system we consider one of the most fundamental and basic problems: counting. We consider the example of the Intelligent Traffic Intersection System (IITS) with a DSPS 302 installed on the centralized server, known as

the traffic control hub. This traffic control hub, for instance, can be a gateway. The traffic control hub uses a DSPS 302 to track the number of cars crossing the intersection in each direction in real time. One of the potential benefits of counting the number of cars is to manage the traffic phase change in real time, i.e. changing the traffic signal from red to green and so on in real time, depending on the traffic flow status at the intersection. Traffic cameras are installed at the intersection, facing each side of the intersection, to capture the video of the surroundings. From a stream analytics context, the video generated by each traffic camera can be considered as a data stream, and each video frame as one data element.

[0019] In this case, let us assume that the data stream processing system gets the following query: count the number of cars crossing the intersection in each direction. The traffic control hub 302 uses video analytics techniques to identify cars in the video stream by extracting useful features from the video in real time. The cars are identified and a count of the cars crossing the intersection from each direction is maintained in real time. In this case, the stream state stored in memory maybe as simple as a single integer variable, say *count*, for each direction of traffic flow. Figure 4 illustrates this problem in a streaming environment. For simplicity and without loss of generality, we show only the count maintained for the cars moving from left to right in the figure.

[0020] For every car that crosses the intersection from left to right, the DSPS 302 at the control hub identifies the car by analyzing the video stream and increases the ‘count’ variable stored in the memory of the DSPS 302 by 1. For instance, in Figure 4A, the control hub 302 observes the first car that has crossed the intersection since the start of analysis. It updates the in-memory state “count” from (initially) 0 to 1. In Figure 4B, the control hub 302 observes a second car crossing the intersection, hence updates the state “count=2”. Similarly, in Figure 4, the state “count” has been updated by the control hub to 5. In order to compute the count of cars, the entire data stream is not required to be stored, i.e., the streaming video as data stream needs to be processed/analyzed in real time. The only information stored in memory for this problem is an updated count, which requires 8 bytes of memory space assuming it is of integer data type. This amount of memory space (8 bytes) is significantly small compared to the gigabytes of streaming video that is generated on a busy intersection.

[0021] The above example is a very simple one. Here, the stream state is directly used to answer the query. But in many complex queries, the stream state generally does not contain the answer directly but is used to compute the answer as and when required. A lot of work is

being done to design efficient algorithms to maintain small space stream state in memory for real time stream analytics even for different problems of varying complexities.

[0022] A distributed data stream processing system is when one or more data streams is processed and analyzed across multiple stream processing nodes to answer a single or multiple queries. Generally, a distributed stream processing system has a coordinator node that coordinates the data communication between the distributed nodes. In the above example (Figure 4), if the control hub with DSPS 302 functionality performs all the stream analytics operations to answer a query, then it is called a centralized stream processing system.

[0023] On the other hand, suppose that some of the video analytics operations are also performed by the traffic cameras instead of the entire analytics being done at the control hub, as shown in Figure 5. Suppose the traffic cameras 1, 2, 3 and 4 capture traffic coming from east, south, west and north respectively. Each of these cameras generates a video data stream. Suppose these cameras are also installed with stream analytics capability. Hence, in the above example of counting vehicles, some of the video analytics operations are performed by the traffic camera, such as identifying the objects that have high probability of being a vehicle. Now, instead of sending the entire video stream, each traffic camera 1, 2, 3 and 4 only send their respective stream state, S1, S2, S3, and S4 respectively, to the control hub. These states comprise only those objects that have a high probability of being a vehicle. The control hub 302 acts as a coordinator here, and receives the state from each of these cameras. Now the control hub 302 can use a union of S1, S2, S3, and S4 to answer the queries such as the following – *Give the overall volume of traffic flow at the intersection*. Since the traffic cameras transmit only their respective stream state to the control hub, instead of the entire stream, the communication cost of the data processing is reduced drastically.

[0024] This kind of a system significantly reduces the communication cost by sending only the necessary information across the communication channel instead of the entire data stream. A distributed data stream processing system also helps in balancing the load of data stream analytics across the distributed system, since now the control hub gets to manage a data stream of significantly lower volume, hence improving the efficiency of the system.

Current Data Stream Processing Platforms

[0025] Work has been done on solving important problems, such as finding maximum/minimum, number of unique elements, frequent items, etc. from a large data stream in real time, using small memory cost. There are many stream processing platforms, open source as well as proprietary, which provide an abstraction to the business or the users to perform data

stream analytics, without getting into the intricacies of the architectural details of data stream processing. Such intricacies may include message queueing, scalability, fault tolerance, parallelism, etc.

[0026] Examples of data stream processing platforms include Apache Storm, IBM Infosphere Stream and Apache Spark Streaming. These stream processing platforms provide means to capture data in a streaming fashion and provide a platform for fault tolerant and reliable stream processing. These platforms may not directly map to the above data stream processing system (shown in Figure 3), which gives a logical view of how a data stream is processed in real time, rather than an architectural view of the platforms supporting stream processing.

[0027] Apache Storm is an open source distributed stream processing platform. It is a reliable, fault tolerant and scalable system. As shown in Figure 6, it has two kinds of nodes: *Master* 602 and *Worker* 604. The Master node 602 runs the Nimbus daemon that assigns tasks to all the worker nodes 604 and monitors failures in the cluster. The Worker node 604 runs the Supervisor daemon that listens for any work assigned by the Nimbus daemon and spawns worker processes to start/stop the work assigned to the worker node. Apache Storm relies on Zookeeper, a centralized service, for the coordination between the Master node 602 and the Worker nodes 604.

[0028] Apache Storm uses three abstractions:

- *Spout* – source of streams within the computation framework of the Storm processing system,
- *Bolt* – comprises of the main computation logic, such as maximum, average, to process the data stream. Each bolt processes a set of input data streams and transforms them into another set of output streams using the computation logic, and
- *Topology* – it is a network of spouts and bolts connected to each other as a “directed acyclic graph”. Every application written in Storm is designed as a topology.

[0029] Figure 7 shows an example of a word count application that counts the occurrence of different words from a stream of sentences. The figure represents a topology of spouts and bolts. Multiple spouts may be on different nodes or a single node of a distributed cluster. The Source sends data stream comprising of sentences (in raw format) across two spouts for load balancing. Each spout in turn forwards the stream, with each sentence as a tuple, to their respective Split bolt. Each Split bolt tokenizes the sentence that it receives and counts the number of occurrences of each word in the sentence. The output stream by each Split bolt, with [word, count] as a tuple, is then forwarded to one of the Count bolts, such that the different Count bolts always receive non-overlapping set of words, to avoid same word being counted at

multiple bolts. Hence, the spouts or bolts within the same vertical column, shown in Figure 7, helps in load balancing and parallelization of application, leading to faster and scalable running of application. There are other ways to implement the word count application and Figure 7 illustrates just one of the approaches.

[0030] IBM Infosphere Streams is another popular distributed stream processing platform. It is a fault-tolerant, reliable and scalable system. All the nodes in the IBM Infosphere Streams distributed cluster may equally participate in executing a streams application. Infosphere Streams has the following main concepts:

- 1) Operators: comprises of computation logic, such as maximum, average, used to transform and manipulate data stream. The operators take a set of data streams as input and transform them based on the computation logic to produce another set of output streams
- 2) Data Flow Graph: Streams application is written in the form of a graph, such that the nodes of the graph are the operators and the edges connecting the operators are the outgoing stream from one operator connecting to another operator as an input stream.
- 3) Processing Element (PE): Individual execution units, comprising of one or more operators, that data flow graph are broken down into to enable distribution and parallelization of streams application.

[0031] Figure 8 shows the runtime execution of IBM Infosphere Streams, where one or more jobs are created corresponding to each runtime instance of IBM Infosphere Streams. Each stream application performing certain analytics operation, such as the word count problem, is submitted to a unique job. This stream application is then divided into more than one PEs based on the data flow graph of the application.

[0032] The concept of writing a data stream application in IBM Infosphere Stream is very similar to Apache Storm. We use the same word count example as used for Apache storm below in Figure 9, to illustrate the running of the streams application in IBM Infosphere Streams. In this example, each operator is executed by individual PEs. The PEs in the same vertical column can be parallelized and the ones across horizontal rows can be distributed across multiple nodes.

[0033] Apache Spark Streaming is another distributed stream processing platform, which has become very popular recently. It is also a fault-tolerant, reliable and scalable system. However, the approach used by Apache Spark Streaming is different as compared to that of the above mentioned platforms. Apache Spark Streaming 1002 uses micro batch processing in order to process data streams, i.e. dividing the data stream into tiny batches and using the batch

processing system, Apache Spark Engine 1004, to process the tiny batches of data, as shown in Figure 10.

[0034] Apache Spark Streaming provides a high level abstraction called discretized stream (D-Stream). D-streams are data streams chopped into batches of small time intervals, called Resilient Distributed Dataset (RDD) and run as small deterministic batch jobs. Figure 11 shows the concept of D-Stream and RDD.

[0035] Figure 12 illustrates the D-stream processing. The data stream is divided into small batches called RDD based on the time intervals. These RDDs go through several rounds of transformation using the computation / aggregation logic to perform the analytics operations, to produce the final output. RDDs in each vertical column can be transformed independent of RDDs in other columns hence enabling parallelization and efficient load distribution.

[0036] The importance of data stream analytics is growing significantly in the IoT/M2M domain. With more and more number of devices connecting to the internet and the growing number of deployments of interconnected IoT/M2M devices, the dependence of these devices and their users on the data is growing too. It is becoming increasingly important to derive some quick insights from these various data streams almost as soon as they are generated. Some of the importance of real time analytics in IoT/M2M are as follows:

- 1) Real time information extraction / decision making: The IoT/M2M sensors and devices generate large amount of data. Enterprises and businesses can gain a lot by analyzing the data, deriving meaningful insights and act on the results quickly. In many IoT/M2M applications, the device/machine requires to make real time decisions autonomously or by using additional context or relationships with other devices to take immediate and appropriate actions.
- 2) Timeliness of Data: In many cases, insights derived from the data generated by a device can result in maximum potential if the data is analyzed in real time, almost as soon as data is generated.
- 3) Limited Storage Resources: Due to limited resources, it can be cost effective if a part of the device data can be discarded after extracting useful information from it. Data stream analytics can filter, aggregate and analyze data without having to store all of it. The storage requirements, in many applications, can be reduced drastically, by processing the data stream in real time and in memory, and only persisting the resulting information which is of value.

- 4) Efficiency: In many cases, even if real-time analysis is not required, it can be much more efficient to process data as it is observed rather than storing the data and then doing batch processing to find information.

[0037] The above points are also applicable for most of the other domains requiring stream analytics, but these requirements can be strongly related with an IoT/M2M system. There are many vendors/companies that provide robust IoT platforms supporting data stream analytics. The following are the main benefits provided by most of these IoT infrastructures, either via the IoT platform or the integrated stream processing system:

- Complex data analytics capabilities
- Fault tolerance
- Distributed service for load balancing and scalability
- Parallel processing

[0038] These features form an integral and important part of adding data stream analytics as a robust service to an IoT/M2M architecture. Below we discuss about a few popular IoT architectures with data stream analytics capabilities. While Microsoft Azure Stream Analytics (Figure 13) and Intel IoT platform (Figures 14A-B) gives a detailed description of the architecture of IoT platforms, Oracle IoT platform (Figure 15) illustrates the architectural framework of the service layer with data stream analytics capabilities. These platforms are very IoT specific and use DSPS concepts discussed above as a component for their data stream analytics.

[0039] Figure 13 shows the architectural framework of the Microsoft Stream Analytics platform. The data stream produced by the Event Producers is received by the Ingestor (the Event Hub), which is based on a publish-subscribe model and is used to consume large volumes of data at high rates. This serves as a single collection point for data streams from various applications. Once a data stream is received by the Ingestor, it is analyzed using the stream processing system and then sent out for long term storage.

[0040] The IoT platform uses Microsoft Azure Streaming service for data stream processing. Microsoft Azure Streaming is a Microsoft proprietary stream processing cloud service used for real time analytics. This system uses complex data stream analytics modules such as a machine learning module. The IoT devices can be registered to this system using the

device registration capabilities. As an alternative, this platform also gives the flexibility of adding Apache Storm as a stream processing system in the infrastructure.

[0041] Figures 14A-B show the detailed architectural framework of the Intel IoT platform. The data stream is transmitted from the Gateway to the stream processing system via a Load Balancer. This system also uses cloud service for data stream analytics.

[0042] Figure 15 shows the IoT service layer proposed by Oracle, which has integrated data stream analytics as one of the services at the service layer, under Event Processing and Big Data & Analytics. The real time analytics service of the Oracle service layer is capable of performing complex event stream processing and uses the query language Continuous Query Language (CQL) for query processing. It is a proprietary service layer for IoT by Oracle, and is made scalable and distributed by using Oracle licensed products.

This section briefly introduces the background information related to the service layer 1602. From a protocol stack perspective, a service layer 1602 is typically situated above the application protocol layer 1606 and provides value added services (e.g. device management, data management, etc.) to applications 1604 (see Figure 16 for illustration) or to another service layer. Hence a service layer 1602 is often categorized as ‘middleware’ services.

[0043] An example deployment of an M2M/IoT service layer 1602, instantiated within a network, is shown in Figure 17. In this example, a service layer instance 1602 is a realization of the service layer 1602. A number of service layer instances 1602 are deployed on various network nodes (i.e. gateways and servers) for providing value-added services to network applications, device applications as well as to the network nodes themselves. Recently, several industry standard bodies (e.g., oneM2M –“oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0”) have been developing M2M/IoT service layers to address the challenges associated with the integration of M2M/IoT types of devices and applications into the deployments such as the Internet, cellular, enterprise, and home networks.

[0044] An M2M service layer 1602 can provide applications and devices access to a collection of M2M-oriented service capabilities. A few examples of such capabilities include security, charging, data management, device management, discovery, provisioning, and connectivity management. These capabilities are made available to applications via Application Programming Interfaces (APIs) which make use of message primitives defined by the M2M service layer 1602.

[0045] The goal of oneM2M is to develop technical specifications which address the need for a common service layer that can be readily embedded within hardware apparatus and

software modules in order to support a wide variety of devices in the field. The oneM2M common service layer supports a set of Common Service Functions (CSFs) (i.e. service capabilities), as shown in Figure 18. An instantiation of a set of one or more particular types of CSFs is referred to as a Common Services Entity (CSE) 1802 which can be hosted on different types of network nodes (e.g. Infrastructure Node (IN) and Middle Node (MN), and Application-Specific Node (ASN)). Such CSEs are termed IN-CSE, MN-CSE and ASN-CSE respectively as defined in oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0. The CSEs provide the service capabilities to other CSEs 1802 as well as to Application Entities (AEs) 1804. Typically, AE 1804 represents an instantiation of application logic for end-to-end M2M solutions and examples of the AE 1804 can be an instance of a fleet tracking application, a remote blood sugar monitoring application, a power metering application, or a controlling application, etc.

[0046] Initially, oneM2M service layer was developed to be compliant to the Resource-Oriented Architecture (ROA) (oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0) design principles, in the sense that different resources are defined within the oneM2M ROA RESTful architecture (as shown in Figure 19). A resource is a uniquely addressable element in the architecture and can be manipulated via RESTful methods such as Create, Retrieve, Update, and Delete. These resources are addressable using Uniform Resource Identifiers (URIs). A resource may contain child resource(s) and attribute(s).

[0047] Recently, oneM2M has started developing an M2M Service Component Architecture (as shown in Figure 20), to consider deployments that are not RESTful based. This architecture is primarily suitable for the infrastructure domain where the CSE is viewed as a set of service components. It largely re-uses the existing service layer architecture shown in Figure 19 but within the service layer it organizes various M2M services and multiple services into service components. In addition to existing reference points, the SoA architecture introduces the inter-service reference point Msc. Communication between M2M Service Components (passing over the Msc reference point) utilizes a web service approach, which is the most popular technology for building Service-Oriented Architecture (SOA)-based software systems.

SUMMARY

[0048] Data Stream analytics (or real time analytics) involves processing and analyzing the streaming data in real time. It is the process of extracting and deriving useful information from continuous and streaming data in real time. A few types of real time analytics operations

that can be performed on a data stream include pattern/anomaly detection, statistical aggregation, predictive analysis, machine learning. Data stream analytics plays a key role in IoT/M2M system due to the need to extract insightful information from the IoT/M2M device generated data in real time.

[0049] Since the IoT service layer acts as a middleware between the M2M devices and the enterprise infrastructure servers, responsible for data management and other services in all the intermediate nodes, the stream analytics capabilities can be integrated at the IoT/M2M service layer. The real time analysis can be performed on the data stream close to the source of generation of data, before the data is stored by the service layers at the data collection points (which is different from the traditional data analytics that are conducted normally after data get stored).

[0050] Existing IoT/M2M platforms depend on cloud services at different levels (such as the cloud services connected to the edge devices or the gateways) of the IoT platforms for data stream analytics. No modular design for physical service layer nodes is described in the literature for data stream analytics. A modular and distributed architecture for data stream processing and analysis is described to incorporate data stream analytics capabilities, called Data Stream Analytics Service (DSAS) in the IoT/M2M service layer. Each service layer node hosting DSAS can be split into two independent modules, Stream Forwarder and Stream Analytics Engine. Stream Forwarder is a light weight processing modules that can be responsible for data preprocessing and routing, whereas Stream Analytics Engine can be responsible for performing actual analytics on the data stream. Separating the two functionalities enables the service layer nodes to efficiently distribute stream analytics tasks across multiple nodes.

[0051] Current service layers, such as oneM2M, lack a standardized service to provide data stream analytics capabilities. Data analytics capabilities are enabled only for the data that has already been stored. A detailed end to end solution of a standard service is described to enable distributed stream analytics capabilities at the IoT/M2M service layer. Having a standard solution for integrating data stream analytics in the service layer will enable data stream real time analysis across multi-domain applications and across different vendors.

[0052] Embodiments include:

- A standard service to enable data stream analytics capabilities at the IoT/M2M service layer, called Data Stream Analytics Service (DSAS)
- Architectural layout and components of DSAS
- Detailed operational procedures of DSAS

- Distributed layout of DSAS hosting nodes in service layer

[0053] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter. Furthermore, the claimed subject matter is not limited to limitations that solve any or all disadvantages noted in any part of this disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0054] A more detailed understanding may be had from the following description, given by way of example in conjunction with accompanying drawings wherein:

[0055] Figure 1 is a diagram that illustrates a data stream in smart home.

[0056] Figure 2 is a diagram that illustrates an exemplary data stream from a GPS probe device.

[0057] Figure 3 is a diagram that illustrates a general logical framework of Data Stream Processing System (DSPS).

[0058] Figure 4A-C are diagrams that illustrate a counting problem in data stream.

[0059] Figure 5 is a diagram that illustrates an exemplary distributed data stream processing system.

[0060] Figure 6 is a diagram that illustrates an Apache Storm architectural layout.

[0061] Figure 7 is a diagram that illustrates an Apache Storm topology for the word count application.

[0062] Figure 8 is a diagram that illustrates IBM Infosphere streams runtime execution.

[0063] Figure 9 is a diagram that illustrates IBM Infosphere streams data flow graph for word count example.

[0064] Figure 10 is a diagram that illustrates micro batch processing in Apache Spark streaming.

[0065] Figure 11 is a diagram that illustrates Apache Spark streaming and Resilient Distributed Dataset (RDD).

[0066] Figure 12 is a diagram that illustrates discretized stream processing in spark streaming.

[0067] Figure 13 is a diagram that illustrates an architectural framework of Azure Stream analytics platform for IoT.

[0068] Figures 14A-B is a diagram that illustrates an architectural framework for IoT Platform by Intel.

[0069] Figure 15 is a diagram that illustrates an architectural framework for IoT Service Layer by Oracle.

[0070] Figure 16 is a diagram that illustrates an exemplary protocol stack supporting a service layer.

[0071] Figure 17 is a diagram that illustrates an exemplary M2M/IoT service layer deployment within a network.

[0072] Figure 18 is a diagram that illustrates Common Services Functions (CSFs) in oneM2M Service Layer.

[0073] Figure 19 is a diagram that illustrates a oneM2M service layer Resource-Oriented Architecture (ROA).

[0074] Figure 20 is a diagram that illustrates a oneM2M Services Component Architecture.

[0075] Figures 21A-B is a diagram that illustrates an Intelligent Intersection Traffic System.

[0076] Figure 22 is a diagram that illustrates a DSAS Layout in an IoT/M2M service layer.

[0077] Figure 23 is a diagram that illustrates details of a DSAS.

[0078] Figure 24 is a flow chart that illustrates an overall procedural description of DSAS.

[0079] Figure 25 is a flow chart that illustrates data source registration.

[0080] Figures 26A-B is a flow chart that illustrates out-of-band stream provisioning.

[0081] Figure 27 is a flow chart that illustrates in-band stream provisioning.

[0082] Figure 28 is a diagram that illustrates data stream parsing.

[0083] Figure 29 is a flow chart that illustrates the procedure for stream parsing.

[0084] Figures 30A-B are a flow chart that illustrates details of query deployment procedure.

[0085] Figure 31 is a diagram that illustrates details of a data stream analytics procedure.

[0086] Figures 32A-B are a flow chart that illustrates details of output triggering.

[0087] Figures 33A-C are diagrams that illustrates distributed layout of DSAS in the service layer.

[0088] Figure 34 is a diagram that illustrates an architecture of the central coordinator in a distributed system.

[0089] Figure 35 is a diagram that illustrates the option of disabling DSAS in the service layer.

[0090] Figure 36 is a diagram that illustrates the option of enabling DSAS in the service layer.

[0091] Figure 37 is a diagram that illustrates a DSAS as a new CSF in oneM2M service layer.

[0092] Figure 38 is a diagram that illustrates the integration of DSAS in oneM2M service layer.

[0093] Figure 39 is a diagram that illustrates a oneM2M resource for DSAS.

[0094] Figure 40 is a diagram that illustrates a `<streamDescriptor>` resource.

[0095] Figure 41 is a diagram that illustrates a `<queryDescriptor>` resource.

[0096] Figure 42 is a diagram that illustrates a `<streamInfo>` resource.

[0097] Figure 43 is a diagram that illustrates an `<output>` resource.

[0098] Figure 44 is a flow chart that illustrates a oneM2M procedure for in-band stream provisioning.

[0099] Figure 45 is a flow chart that illustrates a oneM2M procedure for out-of-band stream provisioning.

[0100] Figures 46A-B are a flow chart that illustrates a oneM2M procedures for query deployment, data stream analytics and output triggering.

[0101] Figure 47 is a diagram that illustrates a GUI for DSAS Clients to trigger a query.

[0102] Figures 48A-B are a diagram that illustrates a GUI for query deployment.

[0103] Figure 49A is a diagram of a M2M/ IoT/WoT communication system that includes a communication network.

[0104] Figure 49B is a diagram of an illustrated M2M service layer in the field domain that provides services for the M2M application, M2M gateway devices, and M2M terminal devices and the communication network.

[0105] Figure 49C is a diagram of an exemplary device that may be used to implement any of the network nodes, devices or apparatuses described herein.

[0106] Figure 49D is a block diagram of a computer system or server that may be used to implement any of the network nodes, devices or apparatuses described herein.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

[00107] In this section, we review the Intelligent Intersection Traffic System (IITS) case with more details in order to justify why service layer needs stream analytics capability and how it can help in this IITS use case or other realistic applications.

[00108] The IITS can be implemented at different levels. At the simplest level, the traffic lights at the intersection are made smart to automatically change phases based on the traffic conditions, and the alert systems are made intelligent to generate real time alerts to warn cars and pedestrians of any possible collisions based on the traffic situation. At the most complex level, the intersection does not have a traffic or stop signal and the vehicles are all expected to be autonomous. A central coordinator at the intersection, say we call it the Intersection Management Gateway, directly communicates with all the vehicles within the communication range from the intersection, guiding them to change lanes, reduce/increase speed or take other necessary actions to cross the intersection efficiently, avoiding collision with other vehicles, bicycles or pedestrians.

[00109] In this system, the end devices, such as the GPS probes and the sensors installed in the vehicles, the traffic cameras at the intersection and other road sensors, transmit continuous stream of traffic data (for instance, GPS data stream by GPS probes, video stream by traffic cameras, etc) to the Intersection Management Gateway, that in turn analyzes the traffic data received from the devices to make real time decisions for efficient traffic management at the intersection. The Intersection Management Gateway generally also forwards this traffic data to the infrastructure servers and cloud servers for further analysis for better future predictions and correlational analysis, and possible permanent storage. In this use case shown in Figures 21A-B, for simplicity, we considered only GPS probe in vehicles and traffic camera as the two types of M2M devices.

[00110] The Intersection Management Gateway is able to perform real time decision making due to the stream analytics services hosted at the gateway. The streaming analytics service at the gateway equips the gateway to process and analyze the traffic data. The stream analytics is mainly performed based on the queries that it receives from devices, users or applications.

[00111] Just for illustration, below are two simple example queries made on this system:

1. *Query 1: Calculate the time to collision between any two vehicles within the communication range of an intersection* – This query uses GPS data and traffic camera video analysis to make real time decisions of how cars should change lanes and change speeds or stop, if need be, to cross the intersection and avoid collisions
2. *Query 2: Calculate the time to collision between any car within communication range of the intersection and a pedestrian* – This query uses GPS data and road sensors to send out real time alerts to pedestrians if there is a risk of collision and warn the cars about the pedestrians if there is a chance of collision.

[00112] The service layer of an IoT/M2M system plays a key role as a middleware, acting as a bridge between the end devices, infrastructure servers and providing necessary services and infrastructure to improve operational efficiencies. The IoT/M2M service layer is responsible for managing data between the end devices and the enterprise infrastructure, along with each intermediate node, via data management and other services, providing needed reliability and security during the entire lifecycle of the data.

[00113] In order to implement the IITS, and similarly other smart IoT/M2M solutions, such that the intersection is efficiently managed in real time, the service layer is given stream analytics capabilities because the data is managed by the service layer during the entire lifecycle of the data. For example, as shown in Figures 21A-B, the Intersection Management Hub receives the data stream continuously over the network channel, and extracts useful traffic information from the data to efficiently manage the traffic at the intersection before the data is forwarded to the storage nodes of the service layer. In this use case, the decision making has to be performed right after the data is generated (i.e. in real time), and before it stored at the intermediate collection points of the service layer. Storage of data into the disk (e.g magnetic disks, Solid State Drives (SSDs)) and later retrieving it for analysis is not suitable for real time analytics for the following reasons:

- Disk Input/Output (I/O) operations due to writing data into or reading data from the disk adds additional significant delays due to I/O overheads. Processing data stored in a disk is slower than processing data on the fly using in-memory information by a magnitude in the order of tens and hundreds.

- The process of storing data into the disk has other overheads too, e.g. indexing the data, which adds further delays unsuitable for real time analysis.

[00114] As stressed before, real time analytics has received a lot of focus by most of the enterprises and industries, in the context of an IoT/M2M system. Hence, it is very important to integrate data stream analytics capabilities in an IoT/M2M system, which has been discussed in the previous use case. Now, since the IoT service layer provides various services for the entire lifecycle of the data within an end to end IoT system, it is very useful to add the data stream analytics capabilities to the service layer of an IoT/M2M system.

[00115] By adding data stream analytics as another service at the IoT/M2M service layer, necessary services can be availed to derive useful insights from the data in real time, and it can be ensured that real time analytics is performed on the data stream before the data is stored at any of the intermediate service layer nodes or permanent storage nodes at the infrastructure servers. There are many studies, such as Fog Computing and Cloudlet, that have proposed techniques to enable real time analytics as close to the data source (edge) as possible. However, the following issues have not been fully addressed by the existing solutions and are the major focus of this work.

[00116] In the current existing IoT platforms, most of the techniques to enable efficient real time analytics in IoT are implemented on cloud servers or virtual machines. The cloud system is built in a hierarchical and federated fashion to optimize data stream analytics near the edge. However, there is almost no focus on building a modular system for incorporating stream analytics capabilities on the physical service layer nodes of IoT/M2M system near the edges, such as on gateways or routers.

[00117] In fact, within IoT/M2M service layer scenario, it is very important to build a modular streaming analytics system at the service layer nodes, with defined independent functionalities for each separate module so that the streaming analytics architecture can be incorporated into service layer with high flexibility and robustness. Back to the example shown in Figures 21A-B, since the Intersection Management Hub is close to the edge, the amount of resources available to the hub might be limited. In such cases, it is desired to have a framework so that complex analytics can be moved to a more powerful node farther away from the edge and light weight stream processing is performed by the nodes near the edge. A modular design of stream analytics capability at service layer nodes can be used.

[00118] Current IoT/M2M service layer standardization efforts do not support data stream analytics - As discussed before, several proprietary as well as open source architectures

have been proposed to integrate stream analytics in IoT. However, existing deployments are mostly proprietary and specific to a certain industry. Currently, there is no standard streaming analytics service in IoT/M2M service layer that can enable streaming analytics capabilities across multi domain applications. Current IoT/M2M service layer standards, such as, oneM2M, provide mechanisms to store the data retrieved from the sensors or other M2M devices to the M2M servers for later retrieval, reasoning and analysis. A service can extract useful information from the data in real time before it is stored anywhere.

[00119] In Figure 22, the Data Stream Analytics Service (DSAS) 2202 at the IoT/M2M service layer integrates streaming analytics capabilities at the layer. In particular, a modular data stream analytics service at the IoT/M2M Service Layer has been defined: DSAS 2202 is designed into two main modules – *a*) a lightweight data stream processing module, called DSAS Stream Forwarder 2212 or DSAS-SF 2212, which performs light weight tasks such as data preprocessing and data routing, and *b*) stream analytics operation based on the user, application or device requirements, which is a more resource consuming data stream processing module, called DSAS Stream Analytics Engine or DSAS-SAE 2210. By separating the two functionalities, we make the data stream analytics architecture more flexible and modular, such that the analytics operation can be easily distributed across the distributed stream analytics system within IoT service layer and cloud.

[00120] In general, DSAS 2202 mainly has the following four main components, DSAS Stream Forwarder (DSAS-SF) 2212, DSAS Stream Analytics Engine (DSAS-SAE) 2210, DSAS Manager 2206, DSAS API 2208.

[00121] DSAS Stream Forwarder (DSAS-SF) 2212 is a light weight stream processing module that

- Acts as an entry point for data stream for stream analytics. Identifies unique data stream with its properties and attributes, and uses access control policy (ACP) to manage control over the data.
- Preprocesses the stream and routes only the reduced stream (only the data stream attributes required for analysis) to the main stream analytics engine (DSAS-SAE 2210), hence controlling the size of the traffic being forwarded to the stream analytics engine
- Acts as a router to forward the incoming stream to the data storage node (may be different from the stream analytics engine), after preprocessing the data as per the requirement of the storage

[00122] DSAS Stream Analytics Engine (DSAS-SAE) 2210 is the module that performs main stream analytics operation such as statistical aggregation, pattern/event detection and predictive analysis. This module receives data stream for analysis from a DSAS-SF 2212 or another DSAS-SAE 2210.

[00123] DSAS Manager 2206 is the main management module of DSAS 2202, which

- Manages the resources allocated to DSAS 2202 such as the physical resources required to store metadata, tables, logs, etc.
- Is responsible for invoking individual services and modules, such as the Security manager 2302, Query manager 2316, etc., within DSAS components 2202 and monitoring these services to make sure they run without errors.
- Is responsible for fault tolerance. In case of failures, it tries to recover the failed processes and jobs in the system, ensuring that the system continues to operate smoothly. One of the approaches is to checkpoint the state of the system periodically.
- Is responsible for providing a secured access over ACP to the external client connecting to the DSAS hosting node to configure and monitor the management services within DSAS 2202 (to be discussed in details under DSAS API 2208 below).
- Is responsible for communicating with other DSAS hosting nodes in a distributed system for fault recovery, load balancing and other communications. In a distributed system, DSAS hosting nodes communicate with each other via their respective DSAS Manager 2206.

[00124] DSAS API 2208 contains the implementation of the Application Programming Interfaces (APIs) for DSAS 2202. The DSAS API 2208 can be used for the following purposes:

- To connect the clients interested in deriving values from the IoT/M2M device data, for instance, users, applications or devices, to the DSAS 2202 hosted on a service layer node, so that the clients can build/deploy queries in DSAS 2202 and access the analytical results outputted by the query deployed in DSAS 2202.
- To manage and control the DSAS 2202 itself – this API 2208 is responsible for configuring the management services, for instance, to
 - manage the fault tolerance policy
 - dynamically configure the access control mechanisms, such as ACP, in the Security manager 2302
 - updating the policy to access resources (e.g. locking mechanism for concurrent reads/writes in a table by multiple entities)

- configure the preprocessor 2306 so that the data stream is preprocessed based on the storage policy of the storage nodes or the information provided by the device guidelines
- update the stream identification policy (in case of Out-of-Band Stream Identification) based on device guidelines.

[00125] Figure 22 shows the general layout of DSAS 2202 in IoT/M2M service layer.

We categorize the nodes connected to the IoT/M2M service layer (SL) into 4 different types:

- 1) DSAS Hosting Service Layer (SL) node 2204: The DSAS 2202 is hosted in these SL nodes. Each DSAS hosting SL node has its own independent DSAS 2202. However, multiple service layer nodes with hosted DSAS 2202 can be connected to each other, via a messaging protocol e.g. MQTT (message queueing telemetry), to provide distributed data stream analytics capabilities to the IoT/M2M service layer.
- 2) Data Source 2218: These nodes are the producers of the IoT/M2M data. The data is transmitted as a stream across the communication channels, starting from these nodes. Examples of a data sources in the previous IITS use case are the traffic cameras generating local traffic video and the GPS probe vehicles generating GPS data. Logically, data source and DSAS hosting SL node are separate entities, but they may be hosted in the same device or physical node.
- 3) DSAS Client node 2216: Hosts the entity, such as a user, an application or an IoT/M2M device, that is interested in deriving useful information or insights from the streaming data in real time. The client 2216 may either be an SL or a non-SL node.
- 4) Data Storage node 2214: An optional node where the data stream is forwarded for storage, if required for later analysis, once useful information has been extracted from it. The storage node 2214 may be a temporary data collection node in the service layer, e.g. oneM2M *<contentInstance>* resource node, or an infrastructure server/data warehouse for permanent storage.

[00126] It is understood that the functionality illustrated in Figure 22, may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, an apparatus of an M2M network (e.g., a server, gateway,

device, or other computer system), such as one of those illustrated in Figures 49C or 49D described below.

[00127] Figure 23 shows the components and modules inside DSAS 2202. As discussed above, DSAS 2202 has four main components – DSAS-SF 2212, DSAS-SAE 2210 and DSAS API 2208 and a main management service to support these components, called DSAS Manager 2206. These components also use shared resources that are stored in the form of tables, trace files, etc. Below, the detailed descriptions of each those components will be presented in each of following sub sections.

[00128] *DSAS Stream Forwarder (DSAS-SF) 2212:* As discussed above, DSAS-SF 2212 is a light weight stream processing component of DSAS 2202. Preprocessing may be done for cleaning data before storage and analysis, discarding unrequired data before being sent to the storage nodes and reducing the dimension of the data stream, per the requirement of DSAS-SAE 2210, before being forwarded to the DSAS-SAE 2210. It may also act as a simple router. The following modules run in DSAS-SF 2212:

- Security Manager 2302 is responsible for secured access of data stream using mechanisms such as an Access Control Policy (ACP), that can be pre-defined by the service layer or vendors or dynamically configured via a management interface. The Security Manager 2302 is mainly used for mutual authentication via access control mechanisms (e.g. ACP). The three main authentications are: 1) ensuring that data stream received from a certain device is authorized to access DSAS 2202, 2) ensuring that a certain data stream from a device is authorized to access DSAS 2202 and 3) ensuring that the DSAS 2202 is authorized to access the incoming data stream of a particular device
- Stream Identification Manager/Parser (SIM/P) 2304 identifies each unique stream, assigns it a unique Stream Identifier if the stream does not already have one, and maintains the table, called Stream ID Store (shown later in Table 2), which stores information of all unique stream received (in other words, observed) by DSAS 2202. It also parses the stream based on its attributes.
- Preprocessor 2306 preprocesses the data generated by the stream. IoT data generally requires preprocessing in order to remove redundancy and noise since from IoT device data are generally very dirty and noisy with several missing points and redundancies. Preprocessing is also used for lossy compression of data using known sampling or aggregation techniques. The preprocessing of data achieves the following main purposes:

1) cleaning the device data, which is generally dirty and full of noise, by removing redundancy and noise from the data, 2) compressing the data by means of aggregation, sampling or other techniques, and 3) the above two purposes in turn reduce the communication cost of the data across transmission channel and the cost of storage of the data.

- Storage Filter 2308 reduces the size of the data stream, forwarded by the preprocessor 2306 for storage, by removing unnecessary attributes from the data stream, based on the policies pre-defined by the storage nodes or the device guidelines, or dynamically configured by a client via the management interface, hence also reducing the transmission and storage cost incurred. The storage policy may also define methods for compressing the data, best optimized for storage.
- Analytics Filter 2310 filters out the attributes from the preprocessed data stream, which are not required for any of the analytics operations, based on the queries deployed in the system. The resulting stream is called the filtered stream. This filtering is done to reduce the dimension of the data stream that is sent by Analytics *Filter* 2310 from DSAS-SF 2212 to *SAE source* 2314 in DSAS-SAE 2210 for actual analytics operation, in order to minimize the size of the data stream handled by DSAS-SAE 2210 for efficiency and reducing the load on DSAS-SAE 2210. In the case of a distributed setup, this also reduces the communication cost if the analytics operation is performed on SAE module of a DSAS 2202 hosted on a different node.

[00129] DSAS Stream Analytics Engine (DSAS-SAE) 2210 can include:

- Query Operators 2312– A query is a mechanism to retrieve specific information from the data stream and/or perform certain actions based on the occurrence of certain conditions or events. These queries pertain to specific data stream analytics operations in the system. Algorithms consisting of computation, statistical, aggregation or more complex logic, are implemented in order to process the query deployed in the system and perform corresponding data stream analysis. The implementations of these algorithms are termed as Query operators 2312. These implementations can be done in native programming languages such as C/C++ or Java, or platform specific languages, depending on the stream analytics engine used within DSAS 2202 to run the query operators 2312. Each query operator maybe used to process one or more queries in DSAS 2202. An instance of

query operator is invoked by the Query manager 2316 to process the corresponding query.

- SAE Source 2314 is the first module in the DSAS-SAE 2210 to receive the filtered data stream from ‘Analytics Filter’ module 2310 in DSAS-SF 2212. The SAE source 2314 refers to the Query ID store 2318 to feed the required stream or set of streams with the desired attributes to each Query Operator instance.
- Query Manager 2316 manages the Query ID store 2318 that contains the metadata and the description of each query deployed in DSAS 2202. The Query Manager 2316 is also responsible for invoking and executing the query operator to process the corresponding query and perform data stream analytics operations. The runtime instance of the query operator, executed by the Query Manager 2316, to process a query is called as Job and the information pertaining to all the jobs, such as the Job ID, corresponding Query ID that the job is processing, job status, corresponding log/trace information, are stored in the Job table 2320 (shown later in Table 9). The Query Manager 2316 is also responsible for the maintenance of the Job table 2320 and monitoring of all the jobs.

[00130] *DSAS Manager 2206:* As discussed above, the primary functions of DSAS Manager 2206 includes resource management and allocation within DSAS hosting node, managing fault tolerance in case of process failures, monitoring management services within DSAS 2202 and in the case of a distributed system, communicating with other DSAS hosting nodes via their respective DSAS Managers 2206.

[00131] *DSAS Application Programming Interface (API 2208):* DSAS API 2208 is an interface built on both the client end and DSAS hosting node end, comprising of a set of routines and protocols that decide how DSAS client 2216 components will interact with DSAS components hosted on the same or different nodes. DSAS API 2208 is used to establish connection between the DSAS hosting SL node and the clients (e.g. IoT devices, applications, users) who are interested in deriving useful insights from the data using DSAS 2202. This connection is qualified by ACP functionality (that can either be pre-defined at the service layer or dynamically configured) to validate client 2216 access into DSAS hosting node.

[00132] The working details of Figure 23 is presented part by part and in details, during the discussion of the detailed procedures below.

[00133] To facilitate the data stream analytics service, some metadata is stored in the form of tables or repositories in the DSAS hosting SL node, such as the Stream ID store (maintains list of unique streams analyzed by DSAS 2202), Query store 3218 (maintains list of queries being processed by DSAS 2202) and Job table 2320 (maintains list of jobs being handled by DSAS 2202 to process the queries). The resources used by these tables are managed by the DSAS manager 2206 and are shared across all the components of DSAS 2202.

[00134] It is understood that the functionality illustrated in Figure 23, may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, an apparatus of an M2M network (e.g., a server, gateway, device, or other computer system), such as one of those illustrated in Figures 49C or 49D described below.

[00135] Figure 24 shows the overall general procedural description of DSAS 2202. The general steps are as follows:

- A. Stream Ingestion – This is the first procedural step in DSAS dealing with data stream. The data stream is ingested into DSAS 2202 via the Security Manager 2302 over ACP, assigned a unique ID for identification purpose, and parsed before being forwarded for further processing. This will be discussed in details later.
- B. Query Deployment – This is required to deploy query in DSAS 2202, based on which the data stream analytics is performed over the data stream. Procedure A and B may occur in any order. A query maybe deployed in DSAS 2202 after the stream ingestion, or the query may have been deployed in DSAS 2202 before even it started receiving the concerned data streams.
- C. Data Stream Analytics – The actual data stream analytics is performed based on the query deployed in the system.
- D. Output Triggering – Though the data stream is processed and analyzed continuously, but the answer to the query may not be outputted continuously. The output maybe explicitly requested by a client 2216 or may be generated as a response to a trigger.
- E. Data Storage – After important information is extracted from the data stream, the data may be forwarded for storage.

[00136] A basic scenario has only one DSAS 2202 in place. As described below, in a distributed scenario, multiple DSAS can be deployed in the system.

[00137] It is understood that the entities performing the steps illustrated in Figure 24 are logical entities that may be implemented in the form of software (i.e., computer-executable

instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 24 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 24. It is also understood that any transmitting and receiving steps illustrated in Figure 24 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

[00138] Procedure A: Stream Ingestion

In general, the Stream Ingestion will have the following three major steps: 1) Data Source Registration, 2) Data Stream Identification and 3) Data Stream Parsing, which will be described in details as follows.

Data Source Registration

[00139] This is the initial set up process for a data source in the IoT/M2M system to register to the service layer node hosting the DSAS 2202. The procedural detail of Data Source Registration has been shown in Figure 25.

[00140] In step 1 of Figure 25, the data source sends a request message to connect (or register) to the DSAS hosting service layer node. The data source may or may not be aware of DSAS 2202 in the service layer node. The data source might simply be interested in registering to the service layer node to forward the data to the next destination.

[00141] In step 2 of Figure 25, the registration of the data stream to DSAS hosting service layer node is performed based on the standard device registration process in the service layer using the standard access control mechanism. However, regardless of the fact that the data source is aware or unaware of the existence of DSAS 2202 in the node, along with the standard device registration, the Security Manager 2302 in DSAS-SF 2212 component also checks the privileges of the data source using an access control mechanism, such as ACP, to ensure that the data source has necessary privileges to avail DSAS 2202. These privileges are specified by the client 2216 (such as users or applications) during the deployment of the data source in the network or later as and when required. The client 2216 uses DSAS API module 2308 to setup privileges for the data sources using the access control mechanism defined for DSAS 2202. The check is performed on the following information of the data source (or device) – the device host

address, the device type, such as a car or a smart phone (if available, since the type of the device may not always be defined) and the device ID (if available, since not all devices may have an ID assigned to it).

[00142] In step 3 of Figure 25, the DSAS hosting SL node sends back a response to the data source confirming the completion of the registration of the data source at the SL node.

[00143] It is understood that the entities performing the steps illustrated in Figure 25 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 25 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 25. It is also understood that any transmitting and receiving steps illustrated in Figure 25 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Data Stream Identification

[00144] This is the process of identifying each unique stream that is received by DSAS 2202 and assign a unique ID, called Unique Identifier, if not already assigned, to each of the unique streams. The data stream identification is done by the Stream Identification Manager / Parser (SIM/P) 2304.

[00145] In particular, it can be done in one of two ways in DSAS 2202: Out-of-Band Stream Provisioning and In-Band Stream Provisioning, which will be described below.

Out-of-Band Stream Provisioning

[00146] Each stream can be uniquely identified in DSAS-SF 2212 by a *Stream Identification Manager*, using the information provided by the client 2216 based on the IoT device guidelines and the pre-provisioning documents. The DSAS 2202 do not depend on the data sources, i.e. the IoT/M2M devices, to obtain stream information for identification. DSAS-SF 2212 does not receive any information from the data source. In this case, the data sources may be completely oblivious about the existence of data stream analytics system in the service

layer. More information about Out-of-Band Stream Provisioning can be found in the procedural description shown in Figures 26A-B.

[00147] In step 1 of Figures 26A-B, the client 2216 (e.g. a user or an application) sends a request via the portal, for instance a GUI or web browser, to establish connection with DSAS 2202. This request contains the client address that will be used for authentication and authorization.

[00148] In step 2 of Figures 26A-B, the client request for connection is received by the DSAS API component 2308 of DSAS 2202, which contains the API 2208 to enable the client side portals to interact with DSAS 2202. This request is forwarded to the DSAS Manager 2206 component.

[00149] In step 3 of Figures 26A-B, the Security manager 2302 within DSAS Manager 2206 uses predefined access control mechanism to check the privileges of the client 2216 within DSAS 2202. It checks whether the client 2216 has the privileges to access DSAS 2202. It also finds the resources and services that the client 2216 is authorized to access. Resources comprise of the tables that a client 2216 can access – read and write access for Stream ID Store 2321 (described later) and Query Store 2318, and read only access for the Log table. Services comprise of the management services that a client 2216 can configure, such as the access control mechanism and the preprocessor configuration.

[00150] In step 4 of Figures 26A-B, the Security manager 2302 in the DSAS Manager 2206 sends back a response to the DSAS API 2208 for the client's connection request. If the client 2216 is authenticated by the Security manager 2302 and has the privileges to access DSAS, the Security manager 2302 also sends the resources and services which the client 2216 is authorized to access.

[00151] In step 5 of Figures 26A-B, the DSAS API 2208 sends the response message to the client portal, e.g. web browser or GUI. If the client 2216 is authenticated, as specified in the response message, then the response message also includes the resources and the services that the client 2216 is authorized to access.

[00152] In step 6 of Figures 26A-B, if the client 2216 has been authenticated by the Security manager 2302 using access control mechanism, as specified by the response message received, then a connection is established between the client 2216 and the DSAS 2202 via the client 2216 side portal, e.g. the web browser or the GUI. The client 2216 gets the view of the resources and services that it can access via the portal. If the client 2216 is authorized to access the Stream ID Store 2321, it enters the information of the data stream it is interested in

provisioning within DSAS 2202, via the portal. Table 1 shows a description of the Stream ID Store 2321. For example, the information entered are stream ID (if available), device information (device address, ID and type, if available) of the device that generates the data stream, stream attributes with metrics for each attribute, and raw data stream format. For instance, consider a raw GPS data stream of ID 'A1' generated from a device of ID 'VH1' and of vehicle type, denoted as 'VH'. The possible information submitted by the client 2216 to DSAS 2202 could be (Normally, in this case, it is assumed that the client 2216 could use the device guidelines or other pre-provisioning document to obtain related information):

- Stream ID: A1
- Device address: IP address of the device and the port number for the port transmitting data
- Device ID: VH1
- Device Type: VH
- List of attributes contained in each message of the data stream: *<Latitude, Longitude, Altitude (meters), Timestamp (hh:mm:ss:ff), Speed (mph)>*
- Message format : comma separated value (CSV)

Table 1: Description of the Stream ID Store 2321

Property name	Multiplicity	Property description
Stream ID	1	A unique identifier for the data stream
Device ID	1	The identifier of the device generating the data stream (the source of the data stream)
Device Type	1	Type of the device generating the data stream, e.g. smartphone
Device Address	1	The host address of the device
Stream Attributes	1..n	List of attributes of the stream, with the respective metrics

Raw Stream Format	1	Format of the raw stream, e.g. CSV, JPEG
Stream Observed	1	No, if the stream has not yet been observed or received by DSAS, else Yes

[00153] In step 7 of Figures 26A-B, the data stream information submitted by the client 2216 through the portal is received by the DSAS API 2208.

[00154] In step 8 of Figures 26A-B, the DSAS API 2208 sends the data stream information submitted by the client 2216 to the DSAS Manager 2206.

[00155] In step 9 of Figures 26A-B, the DSAS Manager 2206 updates the Stream ID Store with all the information provided by the client 2216 that it received via the DSAS API 2208. Table 2 shows the example entries made by the client 2216 for two data streams in Stream ID Store 2321 (still using the IITS use case, i.e., GPS data stream). These streams are only being provisioned and have not yet been observed or received by the DSAS 2202, hence, the last column has been marked as 'No'.

Table 2: Example Entries of the Stream ID Store 2321

Stream ID	Device ID	Device Type	Device Address	List of Stream Attributes with metrics	Raw Stream Format	Stream Observed
A1	VH1	VH	<IP address: port no.>	<Latitude, Longitude, Altitude (meters), Timestamp (hh:mm:ss:ff), Speed (mph)>	CSV	No
A2	VH2	VH	<IP address: port no.>	<Latitude, Longitude, Altitude (meters), Timestamp (hh:mm:ss:ff), Speed (mph)>	Text	No

[00156] In step 10 of Figures 26A-B, the DSAS Manager 2206 sends a confirmation to the DSAS API 2208 about the entry it made in the Stream ID Store with the information provided by the client 2216.

[00157] In step 11 of Figures 26A-B, the DSAS API 2208 sends an acknowledgement to the client 2216, confirming the completion of data stream provisioning, i.e information of a

new unique data stream has been provided to DSAS 2202, to facilitate identification of the corresponding data stream.

[00158] It is understood that the entities performing the steps illustrated in Figures 26A-B are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figures 26A-B may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figures 26A-B. It is also understood that any transmitting and receiving steps illustrated in Figures 26A-B may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

In-Band Stream Provisioning

[00159] In the case of In-band stream provisioning, DSAS-SF 2212 uses the information / metadata directly sent by the data sources to uniquely identify each stream. This information is sent out once a connection is established between the data sources and DSAS-SF 2212, even before the devices start sending out the actual data to the DSAS 2202. In this set up, it is assumed that DSAS 2202 makes itself discoverable on being deployed within the service layer so that the data sources are able to establish a connection with the DSAS 2202. Hence, in this case, the data sources are generally aware of the existence of data stream analytics tool in the system and explicitly send them information about their respective streams. The data stream information is received by the Stream Identification Manager which makes an entry in the Stream ID Store corresponding to the information received for each unique stream. More information about this technique can be found in Figure 27.

[00160] For step 0 of Figure 27, it is considered that the DSAS 2202 has made itself discoverable via the service layer specified service discovery procedure. The data source has been registered to DSAS 2202 using the procedure described in Figure 27. The data source may or may not be aware of the existence of DSAS 2202, during the time of registration, and it may only be interested in registering to the Service Layer node for data forwarding and storage. However, even if the data source were not aware of DSAS 2202 during registration, in some

cases it may become aware of DSAS 2202 later and may explicitly send out its data stream information as a part of Out-of-Band provisioning.

[00161] In step 1 of Figure 27, the data source sends out its data stream information, for which it wants to avail the data stream analytics service, to the DSAS-SF 2212 for stream provisioning, so that DSAS 2202 can uniquely identify the corresponding data stream based on the information provided. The stream information sent by the data source is the same as shown in Step 6 in Figures 26A-B of the procedural description of Out-of-Band Stream Provisioning. Alternatively, if an SL node maintains data stream registry, then the data source only sends out the Stream ID to DSAS 2202, which then performs a lookup to the data stream registry, based on the stream ID provided, to discover other information of the data stream from the semantic description.

[00162] In step 2 of Figure 27, Stream Identification Manager / Parser (SIM/P) 2304 module of DSAS-SF 2212 receives the information from the data source and forwards it to the DSAS Manager 2206, so that corresponding entry is made in the Stream ID Store 2321.

[00163] In step 3 of Figure 27, if the data source has the required privileges to avail DSAS 2202, as checked by the Security manager 2302, an entry for the data source is made in the Stream ID Store 2321 (Table 1), such that the *Device Address* property in the table is set to the host address of the data source, *Device Type* property in the table is set to the type of the device sending out the stream, such as a smartphone, if the information is available, *Device ID* property in the table is set to the device identifier, if available, and the rest of the properties in the table are set to NULL. DSAS Manager 2206 makes an entry in the Stream ID Store 2321 for the data stream information it received from the data source. The example Stream ID Store 2321 entry made by the DSAS Manager 2206 has been shown in Table 2.

[00164] In step 4 of Figure 27, the DSAS Manager 2206 sends a confirmation to the SIM/P module 2304 of DSAS-SF 2212 about the entry it made in the Stream ID Store 2321 with the information provided by the data source.

[00165] In step 5 of Figure 27, SIM/P module 2304 of DSAS-SF 2212 sends an acknowledgement back to the data source, confirming the completion of data stream provisioning, i.e. information of a new unique data stream has been provided to DSAS 2202, to facilitate identification of the corresponding data stream.

[00166] It is understood that the entities performing the steps illustrated in Figure 27 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or

computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 27 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 27. It is also understood that any transmitting and receiving steps illustrated in Figure 27 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Data Stream Parsing

[00167] It is the process of breaking down each stream tuple (or message) of the data stream into individual attributes, so that the attributes are identifiable by the stream processing system. It is also used to retrieve the Stream ID of the concerned data stream for authorization check and used by later analytics operation.

[00168] For instance, consider the following stream tuple <37, 145, 9:41:00.01, 30, A1, VH>. Figure 28 shows output stream tuple on passing input stream tuple through a parser 2304.

[00169] Figure 29 illustrates the stream parsing procedural details to uniquely identify the data stream using the information stored in Stream ID Store 2321.

[00170] In step 1 of Figure 29, data stream is ingested into DSAS 2202 via the SIM/P module 2304 of DSAS-SF 2212 component. The DSAS 2202 may be hosted in the data source itself.

[00171] In step 2 of Figure 29, the streaming data is then passed through the Stream Identification Manager/ Parser (SIM/P) 2304. On receiving the first message of a data stream, SIM/P 2304 matches the data stream with the corresponding entry (or entries) in the Stream ID Store 2321 based on the information of the device from which the data stream is received. If there are multiple streams from a single device, device port number is used to identify the stream. If the device port number is not available, then other information of the stream, such as the data stream format, is used to parse the data stream using a Parser 2304. The parsed data stream attributes are then matched with the Stream ID Store 2321 entries to find the corresponding stream. This step is done to retrieve the stream ID of the data stream.

[00172] In step 3 of Figure 29, if the stream entry is not found in the Stream ID Store 2321 then DSAS 2202 does not proceed with processing the concerned data stream. If the stream

entry of the corresponding stream is found, then the stream ID is used by the Security manager 2302 to check if DSAS 2202 has proper authorization to access the given stream.

[00173] In step 4 of Figure 29, if the data stream clears the authorization check, then the parser 2304 checks if the Stream ID is one of the attributes of the parsed stream. If not, the parser 2304 appends the stream ID as one of the attributes. It is useful for later data stream analytics operation to identify and distinguish between streams. SIM/P 2304, via DSAS Manager 2206, updates the 'Stream Observed' property in Stream ID Store 2321 to 'Yes', if previously set to 'No'.

[00174] In step 5 of Figure 29, the parsed stream is then forwarded to the Preprocessor 2306 for preprocessing of the data stream. In addition to the parsed stream, SIM/P 2304 also forwards the raw data stream to the Preprocessor 2306 for preprocessing before the raw data is sent out for storage. There has been lots of work done on the real time preprocessing of IoT data stream. It is an important part of IoT data processing, since the data is generally dirty and noisy. Dirty data implies that the data may have missing points, or have redundancies. Real time preprocessing methods, called data cleaning, are used to handle missing data and remove noise and redundancies from the data. Data cleaning methods, in some cases, are approximate. Preprocessing is also used to perform compression of the data for reducing the communication cost over the transmission channel. The compression techniques, in some cases, may be lossy. In order to reduce the communication cost, the size of the data may also be reduced using techniques such as sampling, statistical aggregations.

[00175] It is understood that the entities performing the steps illustrated in Figure 29 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 29 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 29. It is also understood that any transmitting and receiving steps illustrated in Figure 29 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Procedure B: Query Deployment

[00176] Query Deployment procedure in DSAS 2202 comprises of the following activities:

[00177] Adding a new query to DSAS 2202 – a DSAS client 2216 adds a new query to DSAS 2202 via DSAS API module 2208 which contains API implementation for the clients to interact with the resources pertaining to the query, such as the Query Store 2318 and the Query operators 2312. Upon adding a new query to DSAS 2202, all the information sent by the client 2216 pertaining to the query are added to the Query Store 2318 (shown in Table 4) and the implementation of the query is saved as a Query Operator. All the properties specified in the Query Store 2318 (described in Table 4) are specified by the client 2216. For instance, adding a query “*Calculate the time to collision between any two vehicles within the communication range of an intersection*” to DSAS 2202 of the service layer of the Intersection Management System in the IITS use case, creates a new entry in the Query Store 2318 as the example shown in Table 5. When a query is added to DSAS 2202, the default value of the Switch property of the corresponding query in the Query Store 2318 can be set as “Enable” or “Disable” depending on the requirement of the application.

[00178] Modifying an existing query in DSAS 2202 – a DSAS client 2216 modifies an existing query. Modifying an existing query implies updating the corresponding query entry in the Query Store 2318 and/or updating the corresponding Query Operator. E.g.: Modify the above query to “*Calculate the time to collision between any two vehicles, and generate an alert if the time to collision goes below 10s*”. The modification of an existing query also includes enabling or disabling a query in the system by updating the value of the *Switch* in the Query Store 2318 for the concerned query.

[00179] Deleting an existing query from DSAS 2202 – a DSAS client 2216 may delete a query from DSAS 2202 via DSAS API, which leads to deleting corresponding entry from the Query Store 2318 and deleting the corresponding Query Operator.

Table 3: Description of the Query Store 2318

Property name	Multiplicity	Property description
Query ID	1	Each query is assigned a unique ID
Query description	1	Description of what query does
Operator name	1	Name of the operator (i.e. the implementation of the algorithm to process the query)
Operator Parameter	1..n	The input parameters accepted by the operator, e.g. window length, i.e. scope of data required for querying
Priority level	1	The priority level of the query to enable DSAS 2202 to rank and prioritize its processing and response time to the query
Input stream ID	1..n	List of input streams on which the query is executed (assuming there are 'm' number of streams in the network)
Attribute list	1..n	The list of attributes, corresponding to each data stream, that is required to process the concerned query
Switch	1	Assumes the value 1 to 'Enable' and 0 to 'Disable' the query in DSAS 2202
Output format	1	Format of the analytical results for the query, e.g. HTML, CSV
Output location	1	Location where the output (analytical results) of the query is stored
Host address	1	Assumes value 'localhost' if <i>Analytics Filter</i> in DSAS-SF forwards the filtered stream to DSAS-SAE in the same host. In a distributed system, DSAS-SF may forward the filtered stream to DSAS-SAE component of another DSAS hosting node for stream analytics, in which case it is the address (IP

address/port number) of that host.

Table 4: Example of a Query Store 2318

Query ID	Query description	Operator name	Operator parameters	Input stream ID	List of attributes for all input streams	Switch (Enable / Disable)
Q1	Continuously compute time to collision between any two vehicles	TTC_V2V	Window length = 1 minute	A1, A2	A1: { <i>Latitude</i> , <i>Longitude</i> , <i>Altitude</i> , <i>Timestamp</i> , <i>Speed</i> } A2: { <i>Latitude</i> , <i>Longitude</i> , <i>Altitude</i> , <i>Timestamp</i> , <i>Speed</i> }	Enable
Q2	Calculate the time to collision between any car within communication	TTC_V2P	Window length = 1 minute	A1, A2	A1: { <i>Latitude</i> , <i>Longitude</i> , <i>Altitude</i> , <i>Timestamp</i> , <i>Speed</i> }	Enable

range of the intersection and a pedestrian					A2: { <i>Latitude</i> , <i>Longitude</i> , <i>Altitude</i> , <i>Timestamp</i> , <i>Speed</i> }	
--	--	--	--	--	--	--

[00180] Figures 30A-B show detailed procedural description of Query Deployment.

[00181] In step 1 of Figures 30A-B, the client 2216 may be a user, an application or an IoT/M2M device who is interested in deriving some useful insight from a data stream observed by DSAS 2202. In order to deploy a query in DSAS 2202, as a first step, the client 2216 initiates a connection with the DSAS hosting node via the DSAS API 2208. The client 2216 sends a request via the portal, for instance a GUI or web browser, to establish connection with DSAS 2202. This request contains the client address that will be used to determine whether the client 2216 has necessary access privileges to connect to DSAS 2202, the resources and services that it has access to, and the kind of access it has for each of the resources.

[00182] In step 2 of Figures 30A-B, the client request for connection is received by the DSAS API 2208 component of DSAS 2202, which contains the API 2208 to enable the client side portals to interact with DSAS 2202. This connection request is forwarded to the DSAS Manager 2206 component, via the management API within DSAS API 2208.

[00183] In step 3 of Figures 30A-B, the Security manager 2302 within DSAS Manager 2206 uses predefined access control mechanism to check the privileges of the client 2216 within DSAS 2202. It checks whether the client 2216 has the privileges to access DSAS 2202. It also finds the resources and services that the client 2216 is authorized to access. Resources comprise of the tables that a client 2216 can access to, such as the Stream ID Store 2321 and the Query Store 2318. The access control mechanism also checks the kind of access the client 2216 have for each of these resources. The client 2216 may have read and write access for Stream ID Store 2321 and Query Store 2318. For the Log table, the client 2216 can only have read access, since the Log table can only be updated by the DSAS modules. More restrained authorization may limit the access of the client 2216 to only specific stream and queries within the Stream ID Store 2321 and Query Store 2318 respectively. Services comprise of the management services that a client 2216 can configure, such as the access control mechanism and the preprocessor configuration.

[00184] In step 4 of Figures 30A-B, the Security manager 2302 in the DSAS Manager 2206 sends back a response to the DSAS API 2208 for the client's connection request. If the client 2216 is authenticated by the Security manager 2302 and has the privileges to access DSAS, the Security manager 2302 also sends the resources and services which the client 2216 is authorized to access (discussed in Step 3 of Figures 30A-B).

[00185] In step 5 of Figures 30A-B, the DSAS API 2208 sends the response message to the client portal, e.g. web browser or GUI. If the client 2216 is authenticated, as specified in the response message, then the response message also includes the resources and the services that the client 2216 is authorized to access.

[00186] In step 6 of Figures 30A-B, if the client 2216 has been authenticated by the Security manager 2302 using access control mechanism, as specified by the response message received, then a connection is established between the client 2216 and the DSAS 2202 via the client 2216 side portal, e.g. the web browser or the GUI. The client gets the view of the resources and services that it can access via the portal. If the client has necessary authorization, as described by the client to perform query deployments, then the client submits required query information via the portal to DSAS 2202.

[00187] Based on the kind of query deployment the client is interested in performing, the client may send out corresponding query information, which may have the following different cases:

- If the client is interested in adding a new query to the system, then the client submits the information, as shown in Table 5, via the portal. The client also deploys an implementation of the query operator within DSAS 2202.

Table 5: Parameters sent by Client for Adding New Query to DSAS 2202

Parameter name	Property description
Query description	Description of what query does
Operator name	Operator implemented to deploy the query
Operator parameters	Set of input parameters to be used while invoking the operator
Input stream ID	List of input streams on which the query is executed

	(assuming there are 'm' number of streams in the network)
Attribute list	The list of attributes, corresponding to each data stream, that is required to process the concerned query
Priority level	Parameter to define the priority level of the query such that the DSAS 2202 can rank and prioritize its processing and response time to the query
Switch	Assumes the value 1 to 'Enable' and 0 to 'Disable' the query in DSAS 2202
Output format	Format of the analytical results for the query, e.g. HTML, CSV
Output location	Location where the output (analytical results) of the query is stored

- If the client is interested in modifying an existing query in DSAS 2202, then the client submits the information, as shown in Table 7, via the portal. The client may be interested in updating just one property of the Query Store 2318. Hence, apart from the Query Description, which the client is required to select to identify the query it is interested in modifying, rest of the parameters which does not require updating by the client is optional. The client may also be interested in updating the query operator implementation, in which case it re-uploads a new query operator implementation or directly updates the implementation within DSAS 2202 via the portal.

Table 6: Parameters sent by Client for Modifying Existing Query to DSAS 2202

Parameter name	Property description
Query description	Description of the query that needs to be modified

Operator name <optional>	Operator implemented to deploy the query
Operator parameters <optional>	Set of input parameters to be used while invoking the operator
Input stream ID <optional>	List of input streams on which the query is executed (assuming there are 'm' number of streams in the network)
Attribute list <optional>	The list of attributes, corresponding to each data stream, that is required to process the concerned query
Priority level <optional>	Parameter to define the priority level of the query such that the DSAS can rank and prioritize its processing and response time to the query
Switch <optional>	Assumes the value 1 to 'Enable' and 0 to 'Disable' the query in DSAS
Output format	Format of the analytical results for the query, e.g. HTML, CSV
Output location	Location where the output (analytical results) of the query is stored

- If the client is interested in deleting a query from DSAS 2202, then the client submits the information, as shown in Table 7, via the portal. Hence, the client only requires to send the Query Description via the portal.

Table 7: Parameters sent by Client for Deleting Existing Query from DSAS

Parameter name	Property description
Query description	Description of the query that needs to be modified

[00188] In step 7 of Figures 30A-B, the query information submitted by the client through the portal is received by the DSAS API 2208.

[00189] In step 8 of Figures 30A-B, the DSAS API 2208 sends the query information, submitted by the client, to the DSAS Manager 2206.

[00190] In step 9 of Figures 30A-B, the DSAS Manager 2206 updates the Query Store 2318 based on the information sent by the client via the portal. In case of adding a new query to the system, the DSAS Manager 2206 assigns a unique Query ID to each new query added to the system, and makes a new entry in the Query Store 2318 against the Query ID, based on the information sent by the client. The Query Operator implementation is also stored in DSAS 2202 via the DSAS Manager 2206. In case of modifying an existing query, the DSAS Manager 2206 refers to the list of parameters sent by the client to update the Query Store 2318 accordingly. The Query Operator is modified via the DSAS Operator, if the client requested for operator implementation update. In case of deleting a query from DSAS 2202, the DSAS Manager 2206 refers to the Query Description sent by the client, and deletes the entry from Query Store 2318 and corresponding operator implementation from DSAS-SAE 2210. The 'Host address' field in Query Store 2318 is set to *localhost* since the same node that receives the query performs stream analytics to answer the query. The value of the 'Host address' in the distributed scenario has been discussed later in the next section (Data Stream Analytics Service as a Distributed Service).

[00191] In step 10 of Figures 30A-B, the DSAS Manager 2206 sends a confirmation to the DSAS API 2208 about updates made to the Query Store 2318 and the query operator.

[00192] In step 11 of Figures 30A-B, DSAS API 2208 sends an acknowledgement to the client, confirming the completion of the query deployment procedure

[00193] Besides the above procedures, some more discussions are presented for Intelligent Query Deployment: As described above, a query can be deployed in DSAS 2202 by an external client, who can create, modify, delete or enable/disable a query in DSAS 2202. However, it is possible to introduce a smart Query manager 2316 in DSAS 2202. A query manager 2316 can be made intelligent by implementing machine learning algorithms in the system, and using the machine learning techniques to train the query manager 2316 to make smart decisions regarding the query.

[00194] A smart query manager 2316 can be used to check for queries, which have not been used for a long time and remove them from the system to free up resources. If required by the applications, it can also be used to make sure that query modifications in DSAS 2202 are not

done frequently. Very frequent query modifications may affect the quality of the results generated and may not be desirable for certain applications. If the DSAS 2202 also has a batch analytics services incorporated in it, then the smart query manager 2316 may also be used identify the appropriate queries which are frequently used in batch analytics, and propose that these be deployed in the streaming analytics scenario, or deploy these queries automatically, if the translation of query from batch analytics to stream analytics is feasible.

[00195] Batch Analytics is a way of processing data, where large volumes of data is collected over a period of time and then, processed and analyzed in a batch. Hence, for batch analytics, data is generally stored in a data repository, such as database or a data warehouse.

[00196] It is understood that the entities performing the steps illustrated in Figures 30A-B are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figures 30A-B may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figures 30A-B. It is also understood that any transmitting and receiving steps illustrated in Figures 30A-B may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Procedure C: Data Stream Analytics

[00197] As mentioned earlier, Data Stream Analytics are performed in order to process all the queries in the Query Store 2318 which have been set to 'Enabled'. Figure 31 shows the detailed workflow for the data stream analytics process.

[00198] In step 1 of Figure 31, as a first step, the query manager 2316 periodically checks for any updates in the Query Store 2318.

- If the Query Store 2318 has been empty before, and the Query manager 2316 just found the first query entry in the Query Store 2318, it reports about the update to DSAS Manager 2206 which invokes the Analytics filter 2310 for the first time, and the preprocessed parsed data stream is now forwarded from Preprocessor 2306 to the

Analytics filter 2310. The Analytics filter 2310 is used to prune out all the attributes from all the data stream observed by DSAS 2202, which are not required by any of the queries stored in the Query Store 2318. The filtered stream to SAE Source 2314 in DSAS-SAE 2210 component. We then move to step 2. The Analytics filter 2310 is implemented in an optimized manner, so that the unnecessary attributes are pruned based on the list of queries and the data stream.

- If the Query Store 2318 is not empty, and there has been no update in the Query Store 2318 and/or Query Operators 2312, then the Query manager 2316 continues to manage the ongoing job.
- If the Query Store 2318 is not empty, and there has been updates in the Query Store 2318 and/or Query Operators 2312, then move to step 2.

[00199] In step 2 of Figure 31, depending on the update in the Query Store 2318, one of the following 4 steps are taken:

- If a new query has been introduced in DSAS 2202 and a new entry is made in the Query Store 2318, then if the 'Switch' entry in the Query Store 2318 is set as enabled, then the Query manager 2316 creates a new job for this query and updates the Job Table 2320 (shown in Table 8). The Job ID field in the table uniquely identifies each job. The Job ID field may be the parent process ID of the job, which can also be used to monitor the related process and child processes. The job invokes the operator that is specified in the Query Store 2318 corresponding to the newly entered query.
- If a query has been modified in DSAS 2202, then if the modification requires re-invocation of the corresponding operator, the Query manager 2316 deletes the previous job from the Job Table 2320 and creates a new job, re-invoking the operator processing the modified query.
- If a query has been deleted in DSAS 2202, then the corresponding job is deleted from the Job Table 2320, killing the corresponding operator instances. The in-memory state are stored in volatile memory so killing the job also deletes the in-memory state automatically.
- If a query switch is changed from disabled to enabled in the Query Store 2318, then the Query manager 2316 creates a new job for this query and updates the Job Table 2320 (shown in Table 8). The job invokes the operator that is specified in the Query Store 2318 corresponding to the newly enabled query. If the query is switched from enabled to disabled, then the corresponding job is deleted by the Query manager 2316 from the Job Table 2320.

Table 8: Description of Job Table 2320

Property name	Multiplicity	Property description
Job ID	1	Each job is assigned a unique ID
Query ID	1	Each job mapped to individual query
Job status	1	Status of the job: running, terminated, running with error, etc
Log Names/Location	1..n	List of names and corresponding location of the log/trace generated by the job

[00200] In step 3 of Figure 31, SAE source 2314 feeds the required set of streams with the desired attributes to each operator corresponding to each query that in turn creates an in-memory state in order to process the query. This in-memory state is maintained continuously.

[00201] It is understood that the entities performing the steps illustrated in Figure 31 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 31 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 31. It is also understood that any transmitting and receiving steps illustrated in Figure 31 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Procedure D: Output Triggering

[00202] We know that the query in a data stream processing system is long standing and continuously running. Though the memory state for a query is maintained continuously, the

state may not be used continuously to answer a query. An answer to the query may be generated continuously or may be required occasionally only when an event is triggered. Output triggering is the step that answers a query based on the requirement of the application. Let us review the Query 1 and Query 2 as discussed in the IITS use case:

- 1) *Query 1: Calculate the time to collision between any two vehicles within the communication range of an intersection*
- 2) *Query 2: Calculate the time to collision between any car within communication range of the intersection and a pedestrian*

[00203] In the above example, output for the Query 1 is triggered continuously because the intersection management hub needs to continuously keep track of the time before collision between any two vehicles close to the intersection. However, output for Query 2 is triggered only when a pedestrian is detected at the intersection, about to cross the intersection. In this case, the event that triggers an output for Query 2 is the detection of a pedestrian at an intersection.

[00204] However a query output may also be triggered manually by a client 3202. For instance, consider a vehicle driver using a smart app to find out the condition of the traffic flow in an intersection. In this case, the smart app may be considered as a client 3202. The driver handling the app may explicitly request for the traffic condition at the intersection. In this example, though the traffic condition is continuously being analyzed in the backend, but it is shown on the app only on being explicitly requested by the driver/app.

[00205] If a client 3202 connects to the DSAS 2202 to trigger a result from the query, its permission is validated over ACP, and then given information about each query by DSAS via DSAS API 2208. The detailed description of Output Triggering procedure is given below.

[00206] If a client 3202 connects to the DSAS 2202 to trigger a query, its permission is validated over ACP, and then given information about each query by DSAS 2202 via DSAS API 2208. The detailed description of Query Triggering procedure is given below.

[00207] In step 1 of Figures 32A-B, the client 3202 may be a user, an application or an IoT/M2M device who is interested in deriving some useful insight from a data stream observed by DSAS 2202. The client 3202 can be a different client 2216 than the client 3202 shown in Figure 22. In order to get an output or results for a query in DSAS 2202, as a first step, the client initiates a connection with the DSAS hosting node via the DSAS API 2208. The client 3202 sends a request via the portal, for instance a GUI or web browser, to establish connection with DSAS 2202. This request contains the client address that will be used to determine whether the

client 3202 has necessary access privileges to connect to DSAS 2202, the resources and services that it has access to, and the kind of access it has for each of the resources.

[00208] In step 2 of Figures 32A-B, the client request for connection is received by the DSAS API 2208 component of DSAS 2202, which contains the API 2208 to enable the client side portals to interact with DSAS 2202. This connection request is forwarded to the DSAS Manager 2206 component, via the management API within DSAS API 2208.

[00209] In step 3 of Figures 32A-B, the Security manager 2302 within DSAS Manager 2206 uses predefined access control mechanism to check the privileges of the client 3202 within DSAS 2202. It checks whether the client 3202 has the privileges to access DSAS 2202. It also finds the resources and services that the client 3202 is authorized to access. Resources comprise of the tables that a client 3202 can access to, such as the Stream ID Store 2321 and the Query Store 2318. The access control mechanism also checks the kind of access the client 3202 have for each of these resources. In this case, the client 3202 requires a read access to the list of query descriptions for which it is interested in getting an output. It also requires access to send necessary parameters for the result to DSAS 2202.

[00210] In step 4 of Figures 32A-B, the Security manager 2302 in the DSAS Manager 2206 sends back a response to the DSAS API 2208 for the client's connection request. If the client 3202 is authenticated by the Security manager 2302 and has the privileges to access DSAS 2202, the Security manager 2302 also sends the resources and services which the client is authorized to access (discussed in Step 3). In this case, the client must receive the list of query description from DSAS 2202 for which it is interested in viewing the result.

[00211] In step 5 of Figures 32A-B, the DSAS API 2208 sends the response message to the client portal, e.g. web browser or GUI. If the client is authenticated, as specified in the response message, then the response message also includes the resources and the services that the client is authorized to access.

[00212] In step 6 of Figures 32A-B, if the client has been authenticated by the Security manager 2302 using access control mechanism, as specified by the response message received, then a connection is established between the client and the DSAS 2202 via the client side portal, e.g. the web browser or the GUI. The client gets the view of the resources and services that it can access via the portal. If the client has necessary authorization, as described by the client to trigger output for specific queries, then the client submits required query information via the portal to DSAS 2202.

[00213] Table 9 shows the parameters that the Client 3202 sends to trigger output for specific queries. The client 3202 can submit request to trigger output for one or more queries

Table 9: Parameters for Triggering Output for Queries within DSAS

Parameter name	Property description
Query description	Description of the query for which output is required
Output Parameters	Set of parameters for the results, e.g. maximum error tolerance, window length, output format, etc

[00214] In step 7 of Figures 32A-B, the query information submitted by the client 3202 through the portal is received by the DSAS API 2208.

[00215] In step 8 of Figures 32A-B, the DSAS API 2208 sends the query information, submitted by the client 3202, to the Query manager 2316 in DSAS-SAE 2210, via DSAS Manager 2206.

[00216] In step 9 of Figures 32A-B, the Query manager 2316 triggers the corresponding running operator instance(s) –the user may be interested in the results for more than one queries

[00217] In step 10 of Figures 32A-B, the operator instances use their corresponding in-memory states that each maintain for their respective query, and the corresponding output parameters sent by the client 3202 for each query, to generate an output for the query.

[00218] In step 11 of Figures 32A-B, the Query manager 2316 sends out the corresponding results of the query to the DSAS API 2208 via DSAS Manager 2206. The output of the query is retrieved by the Query manager 2316 from the output location specified in the Query Store 2318. There may be other ways for the client 2216 to access the output, such as the output being sent by the Query operator 2312 directly to the client 2216 portal via DSAS Manager 2206 and the API 2208.

[00219] In step 12 of Figures 32A-B, DSAS API 2208 sends the required results to the client 3202, as desired by the client 3202.

[00220] It is understood that the entities performing the steps illustrated in Figures 32A-B are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the

method(s) illustrated in Figures 32A-B may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figures 32A-B. It is also understood that any transmitting and receiving steps illustrated in Figures 32A-B may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

Procedure E: Data Storage

[00221] In most of the cases, the data is stored at different data collection points of the service layer for management reasons. In many applications, these data are further forwarded to cloud or infrastructure nodes for warehousing the data for the purpose of later and/or deeper analysis.

Hence, once the SIM/P 2304 module of DSAS has parsed the data stream, the data is passed via the preprocessor 2306 and the Storage filter 2308 for storage, if applicable.

The Storage filter 2308 is used to reduce the dimension of the data stream by pruning out the unnecessary attributes from each tuple, such as attributes which are generated by IoT devices for their own specific usage and may be of no importance in the later part of transmissions. This filter also may use compression techniques, based on the storage policy, to optimize the data for storage. Storage filter 2308 provides the following two benefits to the IoT service layer:

1. Reduces the communication cost of transmitting data across service layer nodes to the infrastructure node.
2. Reduces the storage requirement of the data collection points in the service layer. The storage cost is also reduced drastically for the data warehouses in the cloud or at the infrastructure nodes.

[00222] Storage filter 2308 that reduces the data stream before sending it out for data storage is configured by the client, based on the storage policy of the storage nodes or the device guidelines. These policies define rules that specify which attributes in a data stream and/or the extracted from the pre-processor are stored. For instance, some attributes of a data stream may not be required later, and are present only for the reference of the device generating the data. Using the device guidelines, these attributes can be pruned out from the data stream using the pre-processor.

Data Stream Analytics Service as a Distributed Service

[00223] In this section, we consider the distributed scenario, in which multiple DSAS can be deployed in the system. In particular, a distributed DSAS service is needed for the below reasons:

- Facilitate real time analytics on data stream from distributed data sources. The streaming analytics can be performed on the union of all or subset of the streams from distributed setup in a coordinator node or on cloud using cloud service. For example, the traffic flowing through an individual intersection can leverage real time analytics performed locally at the intersection itself. However, to optimize traffic flow across intersections (e.g. within an entire city), an additional level of analytics can be performed at a central location (e.g. in the cloud). This can be done by individual intersections providing data streams of their results back to the cloud, and the cloud can then perform analytics on these streams and enable intersections to make adjustments in a coordinated manner.
- Facilitate edge analytics – move as much analytics as possible to the edge and move more complex analytics farther from the edge to avoid overloading of the edge nodes
- Load balancing – given massive data, the complex analytics can be distributed across several nodes for faster and possibly parallel analysis
- Scalability – to make the system more scalable with the increasing throughput or size of the data stream

[00224] A distributed set up of DSAS hosting SL nodes will provide the above mentioned advantages to the service layer. Data stream can be communicated to different nodes over light weight message passing models such as MQTT, Extensible Messaging and Presence Protocol (XMPP) etc. These models also ensure reliability on message delivery. All or few of these SL nodes may also be connected to the light weight cloud services, so that the system is robust and may move the analytics operations to cloud as required. A distributed set up of the SL nodes, along with distributed cloud services add flexibility and robustness to DSAS 2202 in the context of data stream analytics.

We have identified the following 3 generic layouts of multiple distributed DSAS nodes in the IoT/M2M service layer.

[00225] Figure 33A shows a distributed setup where DSAS-SF 2212 and DSAS-SAE 2210 components are hosted on two separate DSAS hosting service layer nodes. This covers the case where there are one or more DSAS nodes with DSAS-SF 2212 component, each receiving different set of data streams. These DSAS hosting nodes with DSAS-SF 2212 components are connected to one or more DSAS nodes with DSAS-SAE 2210 component only. This setup is useful specially when the node closer to the edge has low processing capabilities and are only used to preprocess and forward data stream to other nodes for stream analytics and storage.

[00226] The DSAS node with DSAS-SAE 2210 component may be hierarchically connected to DSAS-SAE 2210 components of other DSAS nodes or to the cloud nodes with stream analytics capabilities, to distribute and parallelize the analytics operation across multiple nodes for faster and efficient data analysis.

[00227] Figure 33B shows a scenario where the DSAS hosting node closest to the data source has both DSAS-SF 2212 component for preprocessing and DSAS-SAE 2210 component for analytics operations. This covers the case where there are one or more DSAS nodes with DSAS-SF 2212 component, each receiving different set of data streams. Now a client deploys a query or requests for an output for a query via a DSAS hosting node which is not directly receiving the data stream. Consider that a client has made a request to deploy a query via DSAS # 3 in Figure 33B. Then this DSAS-SAE 2210 will communicate with DSAS-SF#1 and/or DSAS#2 depending on which DSAS-SF 2212 components receive the data streams required for the query. In this case, along with communicating with DSAS-SAE 2210 component of the same nodes for data analytics operations, these DSAS-SF 2212 components will also communicate with one or more DSAS-SAE 2210 component of other DSAS hosting SL nodes. The DSAS node with DSAS-SAE 2210 component may be hierarchically connected to DSAS-SAE 2210 components of other DSAS nodes or to the cloud nodes with stream analytics capabilities, nodes to distribute and parallelize the analytics operation across multiple nodes for faster and efficient data analysis.

[00228] Figure 33C shows a distributed setup where each of the DSAS hosting SL nodes contain both DSAS-SF 2212 and DSAS-SAE 2210 for data preprocessing and stream analytics. In the example given above, where the client makes a request to deploy a query via a DSAS hosting node which is not directly connected to the stream, the DSAS-SAE node of DSAS#3 (Figure 33C) may directly connect to DSAS-SF 2212 components of DSAS#1 and/or

DSAS#2, or may connect via the DSAS-SF 2212 component of its own node. This setup can be useful in a scenario where the first DSAS hosting node performs stream analytics operation but are also connected to other DSAS-SAE nodes and/or cloud with streaming analytics capabilities to for load distribution and parallelization.

[00229] In the figure 33, it is interesting to note that the communication between any two DSAS-SAE 2210 components of different DSAS hosting SL nodes for load distribution and parallelization of the data stream analytics operations mostly depends on the architectural setup and the implementation of the underlying distributed data stream analytics engine. Based on the logic used for distributing the stream analytics operation (for instance, Apache Storm and IBM Infosphere Stream use their own respective logic, Topology and Data Flow Graph respectively, for load distribution and parallelization of data stream application across distributed nodes.)

[00230] However, the communication between DSAS-SF 2212 and DSAS-SAE 2210 follows the same procedural details as specified above for single DSAS hosting node, with the difference that now the components would be hosted in different nodes and communicate with each other via DSAS Manager 2206.

[00231] A distributed system needs to operate in coordination with each other. This coordination is required for communication across distributed system, load balancing of services and streaming analytics operation across distributed system, parallelization of streaming analytics operation of streaming analytics system, scalability, such that the system can be scaled up or down based on the load on the system, and fault recovery, such that in case of failure of services or operations in a node other nodes can resume the failed services or operations. There are two ways to manage coordination between distributed nodes: Centralized Monitoring and Peer-to-Peer monitoring:

[00232] Distributed System with Centralized Monitoring: In the case of centralized monitoring, a single centralized node act as a coordinator for the distributed system. It is responsible for load balancing, fault recovery and communication across distributed system. This setup requires an additional management service in the coordinator node to perform the coordination across distributed node.

[00233] In a distributed system with centralized monitoring, the client connects to the coordinator to deploy query or request for analytical results and the coordinator distributes job across the system. One of the DSAS hosting SL node within the IoT/M2M service layer located on the network or in cloud may act as the coordinator. This node will have an additional

component within DSAS 2202, called Coordinator module. The client communicates with DSAS via this Coordinator module using the DSAS API 2208. This coordinator is used to coordinate communication amongst all the distributed DSAS hosting SL nodes through DSAS Manager 2206 of each node. In this setup, the processes and services running within individual DSAS hosting node are still managed by the DSAS Manager 2206, but in case the DSAS Manager 2206 is unable to recover the failed processes or jobs, it communicates with the coordinator node, so that if possible, the failed job can be resumed in other node(s) by the Coordinator.

[00234] The Host ID field in the Query Store 2318 shown in Table 3 should be the host address of the coordinator node. DSAS-SF 2212 sends out the filtered data stream via the DSAS Manager 2206 to the DSAS-SAE 2210 component of the coordinator node which then deploys it across the distributed system via Coordinator 3402 component.

[00235] It is understood that the functionality illustrated in Figure 33A-C, may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, an apparatus of an M2M network (e.g., a server, gateway, device, or other computer system), such as one of those illustrated in Figures 49C or 49D described below.

[00236] A possible architecture for DSAS with the Coordinator 3402 component has been shown in Figure 34.

[00237] Distributed System with Peer-to-Peer Monitoring: In the case of peer to peer monitoring, each node in the system act as a coordinator 3402 for the distributed system. The client may connect to any of the node to deploy query or request for analytical results and that node is fully responsible for load balancing, fault recovery and communication across distributed system for the query that it receives. In this setup, the DSAS Manager 2206 within the node acts as a coordinator 3402 across the distributed system for the query it received. The Host ID field in the Query Store 2318 shown in Table 3 should be the host address of the DSAS hosting node where the query was deployed. DSAS-SF 2212 sends out the filtered data stream via the DSAS Manager 2206 to the DSAS-SAE 2210 component of the node specified in the host address, which is then responsible for distributing the analytics operation related to the query to other DSAS hosting nodes in the system via DSAS Manager 2206..

[00238] In this setup, the processes and services running within individual DSAS hosting node are still managed by the DSAS Manager 2206, but in case the DSAS Manager 2206 is unable to recover the failed processes or jobs, it communicates with the DSAS Manager 2206

of other DSAS hosting nodes, so that if possible, the failed job can be resumed in other node(s) by the respective DSAS Manager 2206.

Ease of Enabling/Disabling DSAS in Service Layer

[00239] Similar to any other service, DSAS 2202 will require some amount of resources in order to provide data stream analytics capabilities. This overhead might not always be desirable within a service layer. Hence, based on the current requirement of a service layer to process and analyze, it should be flexible to enable or disable DSAS 2202 within a service layer. Figures 35-36 show an equivalent of switch to disable or enable DSAS 2202 in the service layer respectively. Figure 35 shows that if DSAS 2202 is disabled in the service layer, then DSAS hosting node just acts as a router to forward the data to the first data collection point, if located at a different node.

[00240] This section presents the embodiments of the DSAS functionalities within the oneM2M architecture. First, the DSAS 2202 can be integrated as a CSF in the CSE. Then DSAS 2202 related resources and attributes are presented to show the integration with the oneM2M resource tree. Finally, oneM2M procedures are provided to realize the various DSAS procedures.

[00241] It is understood that the functionality illustrated in Figures 35-36, may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, an apparatus of an M2M network (e.g., a server, gateway, device, or other computer system), such as one of those illustrated in Figures 49C or 49D described below.

DSAS as a new CSF

[00242] Common Service Functions (CSF) is defined as informative architectural construct which conceptually groups together a number of sub-functions. Each oneM2M Common Service Entity (CSE) is an instantiation of a set of CSFs of oneM2M environment. A realization of the Data Stream Analytics Service (DSAS) can be as a new CSF to the oneM2M CSE (DSAS CSF 3702) as shown in Figure 37. A DSAS hosting CSE 3704 can include the DSAS CSF 3702.

Architectural Layout of DSAS in oneM2M Service Layer

[00243] Data Stream Analytics Service (DSAS) 2202 can be integrated into the oneM2M service layer architecture as shown in Figure 38. The Application Entity (AE), comprising of sensors and/or other IoT/M2M devices, can be considered as the source of the data stream or the data sources. There may be one or more DSAS hosting CSEs 3704. Multiple DSAS hosting CSEs 3704 imply a distributed set up of the data stream analytics service in the oneM2M service layer. The DSAS hosting CSE 3704 may be an ASN-CSE, an MN-CSE, or an IN-CSE, whether the CSE can be a physical server or a cloud server. The data storage CSE 3802 may be an MN-CSE that temporarily collects data and stores it as <contentInstance> resources, or it may be an IN-CSE where the data is permanently stored. The data storage node may also be a oneM2M cloud server.

[00244] In this architecture, DSAS 2202 is used to analyze a data stream before even the data is stored for the first time as a <contentInstance> resource, for faster and real time analysis of data. Hence, DSAS 2202 is hosted in a CSE that lies in between an AE and the data storage CSE. The DSAS 2202 may be hosted within the data storage CSE as well. In that case, data stream analytics is performed before data is stored in the storage node. A DSAS client 2216 interested in deriving useful insights from the data stream may be an AE or a CSE from the same network or a different network.

[00245] It is understood that the functionality illustrated in Figures 37-38, may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, an apparatus of an M2M network (e.g., a server, gateway, device, or other computer system), such as one of those illustrated in Figures 49C or 49D described below.

oneM2M Resource for DSAS

[00246] A new oneM2M resource type for DSAS 2202, called <DSAS>, is described within the oneM2M system as shown in Figure 39. This <DSAS> resource can be centrally located at the CSE base where it hosts all DSAS 2202 related resources. Alternatively, it may also exist as a child resource of an <AE>, a <node>, or other oneM2M resources. The resource contains an attribute *switch* and two child resources, <streamDescriptor> and <queryDescriptor>. The new attribute is shown in Table 10 while the child resources of <DSAS> are shown in Table 11.

[00247] The attribute *switch* shall be used to enable or disable DSAS 2202 within the CSE and can assume two values – 0 or 1. The value 0 disables DSAS 2202 in the oneM2M service layer as shown in Figure 35 and the value 1 enables DSAS 2202 in the oneM2M service layer as shown in Figure 36.

Table 10: Attributes of <DSAS>

Attributes	Multiplicity	Description
<i>switch</i>	1	To enable or disable DSAS within a CSE. Can assume value 0 or 1 to enable or disable DSAS within a CSE, respectively.

[00248] The resource <*streamDescriptor*> shall describe each unique data stream received by the DSAS 2202. It contains the metadata and additional information of the stream it is describing. It is the mapping of the *Stream ID Store 2321* shown in Table 2 to the oneM2M system. The resource <*queryDescriptor*> shall be used to store all the information pertaining to the queries built into the DSAS 2202, one resource per query. It is the mapping of the *Query Store 2318* shown in Table 3 to the oneM2M system.

Table 11: Child resources of <DSAS>

Child Resources	Multiplicity	Description
< <i>streamDescriptor</i> >	0..n	Contains information regarding each unique stream received by DSAS 2202.
< <i>queryDescriptor</i> >	0..n	Contains information regarding all the queries built in DSAS 2202.

The <*streamDescriptor*> Resource

[00249] The resource tree for the <*streamDescriptor*> resource is shown in Figure 40. Each stream the DSAS 2202 monitors will have its own <*streamDescriptor*> resource, which could be created via In-Band or Out-of-Band methods.

[00250] The list of attributes of the *<streamDescriptor>* resource are described in Table 12.

Table 12: Attributes of *<streamDescriptor>*

Attributes	Multiplicity	Description
<i>streamID</i>	1	Contains a unique identifier for the stream that the resource <i><streamDescriptor></i> is describing.
<i>AEID</i>	0..1	The identifier of the AE (device) that generates the data stream being described, only if the device identifier exists.
<i>AEType</i>	0..1	The type of the AE (device) that generates the data, if only the AE type exists, e.g smartphone
<i>AEAddress</i>	1	The host address of AE, containing the IP address and the port number
<i>rawStreamFormat</i>	1	Format of the raw data stream generated by the AE, e.g CSV, jpeg
<i>tupleFormat</i>	1..n	Contains the list of tuple attributes of the stream being described.
<i>dataStoreURI</i>	0..1	The path or URI to the <i><container></i> resource where the pre-processor output data stream is stored, after DSAS has extracted necessary information from the data stream. The stored data consists of <i><contentInstance></i> resources that are created.
<i>dataOutputSelect</i>	1	This indicator selects the type of data to be saved in the Data Store node after passing through the pre-processor. Output selection can be: – pre-processed, selects pre-process data to be stored – raw, selects raw data to be stored – none, don't select any data to be stored

[00251] The child resource of the *<streamDescriptor>* resource is described in Table 13.

Table 13: Child resources of <streamDescriptor>

Child Resources	Multiplicity	Description
<subscription>	0..n	To subscribe to the <streamDescriptor> resource. Refer to oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0 for details.

The <queryDescriptor> Resource

[00252] The resource tree for the <queryDescriptor> resource is shown in Figure 41. The <queryDescriptor> resource provides the query operation that is used by the data stream analytic engine.

[00253] The list of attributes of the <queryDescriptor> resource are described in Table 14.

Table 14: Attributes of <queryDescriptor>

Attributes	Multiplicity	Description
<i>queryID</i>	1	Contains a unique ID for the query being described by the <queryDescriptor>, distinguishing it from other query in DSAS hosting CSE
<i>queryDescription</i>	1	Description of the analytical operation pertaining to the query
<i>operatorName</i>	1	Name of the operator (query implementation), that has been deployed in the DSAS hosting CSE
<i>operatorParameter</i>	1..n	The input parameters accepted by the query operator, e.g. window length, i.e. scope of data required for querying
<i>priorityLevel</i>	1	The priority level of the query to enable DSAS hosting to rank and prioritize its processing and response time to the query
<i>hostAddress</i>	1	Assumes value 'localhost' if <i>Analytics Filter</i> in DSAS-SF forwards the filtered stream to DSAS-SAE in the same host. In a distributed system, DSAS-SF may forward the filtered stream to DSAS-SAE component of another DSAS hosting node for stream analytics, in which case it is the address (IP address/port number) of that host.
<i>switch</i>	1	Assumes values 0 or 1 to disable or enable a query invocation within CSE, respectively

[00254] The child resource of the <queryDescriptor> resource is described in Table 15.

Table 15: Child resources of <queryDescriptor>

Child Resources	Multiplicity	Description
<streamInfo>	1..n	Information of a set of data streams which the query described by the corresponding <queryDescriptor> is interested in performing analytics upon
<output>	0..n	Information regarding the output generated by the query, described by the corresponding <queryDescriptor>
<subscription>	0..n	To subscribe to the output of the concerned query, described by corresponding <queryDescriptor>. Refer to oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0 for details.

The <streamInfo> Resource

[00255] The resource tree of the <streamInfo> resource is shown in Figure 42. The <streamInfo> resource provides the data source input(s) used by the data stream analytics engine to perform the required query.

[00256] The attributes of the <streamInfo> resource are described in Table 16.

Table 16: Attributes of <streamInfo>

Attributes	Multiplicity	Description
<i>streamID</i>	1	Unique identifier of the data stream on which the query, being described by the corresponding <queryDescriptor>, is made
<i>tupleFormat</i>	1..n	List of attributes of the concerned data stream, which is required to process and answer the query, described by corresponding <queryDescriptor>

1.1.1.1 The <output> Resource

[00257] The resource tree of the <streamInfo> resource is shown in Figure 43. The information in the <output> resource directs the DSAS 2202 on the format of the output that is triggered by the corresponding query and where to store the output.

[00258] The attributes of the <output> resource are described in Table 17.

Table 17: Attributes of <output>

Attributes	Multiplicity	Description
<i>outputName</i>	0..1	Contains the name of the output generated by the concerned query, if applicable, based on the corresponding analytics operation performed
<i>outputURI</i>	0..1	Contains the Uniform Resource Identifier (URI) for the location where the output generated by the concerned query is stored
<i>outputFormat</i>	0..1	Contains the format in which the output given by the concerned query is generated, e.g. CSV, HTML

[00259] The child resource of the <output> resource is described in Table 18.

Table 18: Child resources of <streamDescriptor>

Child Resources	Multiplicity	Description
<subscription>	0..n	To subscribe to the <output> resource of the concerned query. Refer to oneM2M-TS-0001 oneM2M Functional Architecture-V-2.4.0 for details.

[00260] oneM2M Procedures for Data Stream Identification and Provisioning

Figure 44 shows the procedure for In-Band Stream Provisioning within the oneM2M system. In this case, DSAS 2202 is integrated within the CSE (DSAS hosting CSE 3704).

[00261] In step 0 of Figure 44, it is considered that all the configuration pertaining to the oneM2M system is performed so that the entire DSAS architecture works together. The data source AE 4402, CSE 3704, and DSAS 2202 are all able to communicate with each other.

[00262] In step 1 of Figure 44, a request is sent by the AE (device generating data stream) to the CSE 3704 to create a <streamDescriptor> resource for the data stream transmitted by the AE. The AE includes information and metadata of the data stream in the request.

[00263] In step 2 of Figure 44, the CSE 3704 checks the ACP to ensure that AE 4402 has access rights to create the <streamDescriptor> resource. If access control is granted, continue on to Step 3; otherwise, go to Step 4.

[00264] In step 3 of Figure 44, the CSE 3704 creates the <streamDescriptor> resource based on the information received from the AE. An entry for the Stream ID Store 2321 within DSAS 2202 is also created.

[00265] In step 4 of Figure 44, CSE 3704 sends a response with appropriate status to the AE.

[00266] It is understood that the entities performing the steps illustrated in Figure 44 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 44 may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when

executed by a processor of the apparatus, perform the steps illustrated in Figure 44. It is also understood that any transmitting and receiving steps illustrated in Figure 44 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

[00267] Figure 45 shows an embodiment of Out-of-Band Stream Provisioning within the oneM2M system. DSAS 2202 is shown external to the CSE 4504 but it may also be integrated within the CSE 4504. Note that AE(DSAS Client) 4502 could be another CSE acting as a DSAS Client.

[00268] In step 0 of Figure 45, it is considered that all the configurations pertaining to the oneM2M system is performed so that the entire DSAS architecture works together. DSAS 2202, CSE, and DSAS Client AE 4502 are all able to communicate with each other.

[00269] In step 1 of Figure 45, a request is sent by AE (DSAS Client) to the CSE 4504 to create a <streamDescriptor> resource for the data stream transmitted by another AE (not shown). AE includes information and metadata of the data stream in the request.

[00270] In step 2 of Figure 45, the CSE checks the ACP to ensure that DSAS Client AE 4502 has access rights to create the <streamDescriptor> resource. If access control is granted, continue on to Step 3; otherwise, go to Step 7.

[00271] In step 3 of Figure 45, the CSE 4504 creates the <streamDescriptor> resource based on the information received from the DSAS Client AE 4502.

[00272] In step 4 of Figure 45, CSE 4504 sends a request to DSAS to create an entry in the Stream ID Store 2321. For the case DSAS is integrated within the CSE 4504, this step is an internal process.

[00273] In step 5 of Figure 45, DSAS creates entry in the Stream ID Store 2321.

[00274] In step 6 of Figure 45, DSAS sends response with status of Stream ID Store 2321 entry creation. For the case DSAS is integrated within the CSE 4504, this step is an internal process.

[00275] In step 7 of Figure 45, CSE 4504 sends a response with appropriate status to the DSAS Client AE 4502.

[00276] It is understood that the entities performing the steps illustrated in Figure 45 are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figure 45 may be implemented in the form of software (i.e., computer-executable

instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figure 45. It is also understood that any transmitting and receiving steps illustrated in Figure 45 may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

oneM2M Procedures for Query Deployment, Data Stream Analytics and Output Triggering

[00277] Figures 46A-B shows an embodiment of oneM2M procedures for Query Deployment, Data Stream Analytics, and Output Triggering. DSAS is shown located separately from the CSE 4504 but it may also be integrated within the CSE 4504. When integrated, Steps 2 to 7 run within the CSE 4504 and messaging occurs internally between CSE 4504 and DSAS 2202. Note that either AE2(DSAS Client) 4502 or AE3(DSAS Client) 4602 could be another CSE acting as a DSAS Client.

[00278] In step 0 of Figures 46A-B, it is considered that all the configurations pertaining to the system are performed so that the entire DSAS architecture works together. The data stream is ingested from the data source to the DSAS 2202 and is a continuous stream. The stream may be ingested before or after the query is deployed in the system.

[00279] In step 1 of Figures 46A-B, AE2 (DSAS Client) 4502 performs a query deployment by sending a request to create a new <queryDescriptor> resource in the CSE 4504.

[00280] In step 2 of Figures 46A-B, the CSE 4504 checks ACP for access control, and if granted, creates a new <queryDescriptor> resource using the information provided by AE24502.

[00281] In step 3 of Figures 46A-B, the CSE 4504 informs DSAS 2202 regarding the new <queryDescriptor> resource. For the case DSAS 2202 is integrated within the CSE 4504, this step is an internal process.

[00282] In step 4 of Figures 46A-B, the DSAS uses the information sent by the CSE 4504 to deploy the corresponding operator implementation of the query within DSAS-SAE 2210. The Query manager 2316 invokes this operator, if the query is set as “Enabled” (the attribute *switch* should have value ‘1’) within <queryDescriptor> resource. The data stream analytics is started if query deployment is successful.

[00283] In step 5 of Figures 46A-B, DSAS 2202 sends a confirmation to the CSE 4504 regarding the completion of the query deployment. For the case DSAS 2202 is integrated within the CSE 4504, this step is an internal process.

[00284] In step 6 of Figures 46A-B, the CSE 4504 sends an acknowledgement to AE2 regarding the completion of the query deployment.

[00285] In step 7 of Figures 46A-B, the data stream analytic processing continues. This processing was started in Step 4 and this step represents the continuous operation of the analytic processing as long as the query deployment is successful.

[00286] In step 8 of Figures 46A-B, AE3 (DSAS Client) 4602, which may be the same AE as the one that deployed query or a different AE, is interested in obtaining results from the deployed query. It sends a subscription request to CSE 4504 to subscribe to the <output> resource of <queryDescriptor>.

[00287] In step 9 of Figures 46A-B, the CSE 4504 completes the subscription process for AE3 and sends back a confirmation regarding the same.

[00288] In step 10 of Figures 46A-B, the output for the subscribed query is triggered at the DSAS based on some triggered event or externally by a DSAS client.

[00289] In step 11 of Figures 46A-B, DSAS 2202 sends the output to the CSE 4504 which stores it in the <output> resource.

[00290] In step 12 of Figures 46A-B, CSE 4504 then sends the output to AE3 4602.

[00291] In step 13 of Figures 46A-B, AE3 4602 send a confirmation for receiving the output.

[00292] It is understood that the entities performing the steps illustrated in Figures 46A-B are logical entities that may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of, and executing on a processor of, a network apparatus or computer system such as those illustrated in Figure 49C or Figure 49D. That is, the method(s) illustrated in Figures 46A-B may be implemented in the form of software (i.e., computer-executable instructions) stored in a memory of a network apparatus, such as the apparatus or computer system illustrated in Figure 49C or Figure 49D, which computer executable instructions, when executed by a processor of the apparatus, perform the steps illustrated in Figures 46A-B. It is also understood that any transmitting and receiving steps illustrated in Figures 46A-B may be performed by communication circuitry of the apparatus under control of the processor of the apparatus and the computer-executable instructions (e.g., software) that it executes.

[00293] Interfaces, such as Graphical User Interfaces (GUIs), can be used to assist user to control and/or configure functionalities related to data stream analytics in the service layer. Figure 47 shows a GUI 4702 for the client to create queries or to access the queries that are deployed and enabled in DSAS 2202. The client is given the list of queries enabled in the system via a dropdown menu. Also, the client is given an option to select the format in which the output of the query is saved. The output may be saved as a CSV file, text file, HTML, etc. If the query supports timed window, then the client also has an option of selecting the window length, i.e. the scope of the data for which the query was processed. Since data stream analytics may involve approximate answers, the client is also given an option of selecting the maximum error it can tolerate in its answer.

[00294] The GUI to deploy a query can be interactive, where based on the kind of query deployment the client is interested in performing (assuming the client has necessary authorization), the GUI gives different options to the client. Exemplary interfaces provided to the client based on the kind of deployment they are interested in performing are shown in Figures 48A-B.

[00295] It is to be understood that interfaces such as that of Figures 47 and 48 can be produced using displays such as those shown in Figures 49C-D described below.

Example M2M/IoT/WoT Communication System

[00296] The various techniques described herein may be implemented in connection with hardware, firmware, software or, where appropriate, combinations thereof. Such hardware, firmware, and software may reside in apparatuses located at various nodes of a communication network. The apparatuses may operate singly or in combination with each other to effect the methods described herein. As used herein, the terms “apparatus,” “network apparatus,” “node,” “device,” and “network node” may be used interchangeably.

[00297] The service layer may be a functional layer within a network service architecture. Service layers are typically situated above the application protocol layer such as HTTP, CoAP or MQTT and provide value added services to client applications. The service layer also provides an interface to core networks at a lower resource layer, such as for example, a control layer and transport/access layer. The service layer supports multiple categories of (service) capabilities or functionalities including a service definition, service runtime enablement, policy management, access control, and service clustering. Recently, several industry standards bodies, e.g., oneM2M, have been developing M2M service layers to address

the challenges associated with the integration of M2M types of devices and applications into deployments such as the Internet/Web, cellular, enterprise, and home networks. A M2M service layer can provide applications and/or various devices with access to a collection of or a set of the above mentioned capabilities or functionalities, supported by the service layer, which can be referred to as a CSE or SCL. A few examples include but are not limited to security, charging, data management, device management, discovery, provisioning, and connectivity management which can be commonly used by various applications. These capabilities or functionalities are made available to such various applications via APIs which make use of message formats, resource structures and resource representations defined by the M2M service layer. The CSE or SCL is a functional entity that may be implemented by hardware and/or software and that provides (service) capabilities or functionalities exposed to various applications and/or devices (i.e., functional interfaces between such functional entities) in order for them to use such capabilities or functionalities.

[00298] Figure 49A is a diagram of an example machine-to machine (M2M), Internet of Things (IoT), or Web of Things (WoT) communication system 10 in which one or more disclosed embodiments may be implemented. Generally, M2M technologies provide building blocks for the IoT/WoT, and any M2M device, M2M gateway, M2M server, or M2M service platform may be a component or node of the IoT/WoT as well as an IoT/WoT service layer, etc. Communication system 10 can be used to implement functionality of the disclosed embodiments and can include functionality and logical entities such as such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48.

[00299] As shown in Figure 49A, the M2M/ IoT/WoT communication system 10 includes a communication network 12. The communication network 12 may be a fixed network (e.g., Ethernet, Fiber, ISDN, PLC, or the like) or a wireless network (e.g., WLAN, cellular, or the like) or a network of heterogeneous networks. For example, the communication network 12 may be comprised of multiple access networks that provide content such as voice, data, video, messaging, broadcast, or the like to multiple users. For example, the communication network 12

may employ one or more channel access methods, such as code division multiple access (CDMA), time division multiple access (TDMA), frequency division multiple access (FDMA), orthogonal FDMA (OFDMA), single-carrier FDMA (SC-FDMA), and the like. Further, the communication network 12 may comprise other networks such as a core network, the Internet, a sensor network, an industrial control network, a personal area network, a fused personal network, a satellite network, a home network, or an enterprise network for example.

[00300] As shown in Figure 49A, the M2M/ IoT/WoT communication system 10 may include the Infrastructure Domain and the Field Domain. The Infrastructure Domain refers to the network side of the end-to-end M2M deployment, and the Field Domain refers to the area networks, usually behind an M2M gateway. The Field Domain and Infrastructure Domain may both comprise a variety of different network nodes (e.g., servers, gateways, device, and the like). For example, the Field Domain may include M2M gateways 14 and terminal devices 18. It will be appreciated that any number of M2M gateway devices 14 and M2M terminal devices 18 may be included in the M2M/ IoT/WoT communication system 10 as desired. Each of the M2M gateway devices 14 and M2M terminal devices 18 are configured to transmit and receive signals, using communications circuitry, via the communication network 12 or direct radio link. A M2M gateway 14 allows wireless M2M devices (e.g. cellular and non-cellular) as well as fixed network M2M devices (e.g., PLC) to communicate either through operator networks, such as the communication network 12 or direct radio link. For example, the M2M terminal devices 18 may collect data and send the data, via the communication network 12 or direct radio link, to an M2M application 20 or other M2M devices 18. The M2M terminal devices 18 may also receive data from the M2M application 20 or an M2M terminal device 18. Further, data and signals may be sent to and received from the M2M application 20 via an M2M service layer 22, as described below. M2M terminal devices 18 and gateways 14 may communicate via various networks including, cellular, WLAN, WPAN (e.g., Zigbee, 6LoWPAN, Bluetooth), direct radio link, and wireline for example.

[00301] Exemplary M2M terminal devices 18 include, but are not limited to, tablets, smart phones, medical devices, temperature and weather monitors, connected cars, smart meters, game consoles, personal digital assistants, health and fitness monitors, lights, thermostats, appliances, garage doors and other actuator-based devices, security devices, and smart outlets.

[00302] Referring to Figure 49B, the illustrated M2M service layer 22 in the field domain provides services for the M2M application 20, M2M gateway devices 14, and M2M terminal devices 18 and the communication network 12. Communication network 12 can be

used to implement functionality of the disclosed embodiments and can include functionality and logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48. The M2M service layer 22 may be implemented by one or more servers, computers, devices, virtual machines (e.g. cloud/storage farms, etc.) or the like, including for example the devices illustrated in figures 49C and 49D described below. It will be understood that the M2M service layer 22 may communicate with any number of M2M applications, M2M gateways 14, M2M terminal devices 18, and communication networks 12 as desired. The M2M service layer 22 may be implemented by one or more nodes of the network, which may comprises servers, computers, devices, or the like. The M2M service layer 22 provides service capabilities that apply to M2M terminal devices 18, M2M gateways 14, and M2M applications 20. The functions of the M2M service layer 22 may be implemented in a variety of ways, for example as a web server, in the cellular core network, in the cloud, etc.

[00303] Similar to the illustrated M2M service layer 22, there is the M2M service layer 22' in the Infrastructure Domain. M2M service layer 22' provides services for the M2M application 20' and the underlying communication network 12 in the infrastructure domain. M2M service layer 22' also provides services for the M2M gateways 14 and M2M terminal devices 18 in the field domain. It will be understood that the M2M service layer 22' may communicate with any number of M2M applications, M2M gateways and M2M devices. The M2M service layer 22' may interact with a service layer by a different service provider. The M2M service layer 22' by one or more nodes of the network, which may comprises servers, computers, devices, virtual machines (e.g., cloud computing/storage farms, etc.) or the like.

[00304] Referring also to Figure 49B, the M2M service layers 22 and 22' provide a core set of service delivery capabilities that diverse applications and verticals can leverage. These service capabilities enable M2M applications 20 and 20' to interact with devices and perform functions such as data collection, data analysis, device management, security, billing, service/device discovery etc. Essentially, these service capabilities free the applications of the burden of implementing these functionalities, thus simplifying application development and

reducing cost and time to market. The service layers 22 and 22' also enable M2M applications 20 and 20' to communicate through networks 12 in connection with the services that the service layers 22 and 22' provide.

[00305] The methods of the present application may be implemented as part of a service layer 22 and 22'. The service layer 22 and 22' is a software middleware layer that supports value-added service capabilities through a set of Application Programming Interfaces (APIs) and underlying networking interfaces. Both ETSI M2M and oneM2M use a service layer that may contain the connection methods of the present application. ETSI M2M's service layer is referred to as the Service Capability Layer (SCL). The SCL may be implemented within an M2M device (where it is referred to as a device SCL (DSCL)), a gateway (where it is referred to as a gateway SCL (GSCL)) and/or a network node (where it is referred to as a network SCL (NSCL)). The oneM2M service layer supports a set of Common Service Functions (CSFs) (i.e. service capabilities). An instantiation of a set of one or more particular types of CSFs is referred to as a Common Services Entity (CSE) which can be hosted on different types of network nodes (e.g. infrastructure node, middle node, application-specific node). Further, connection methods of the present application can be implemented as part of an M2M network that uses a Service Oriented Architecture (SOA) and/or a resource-oriented architecture (ROA) to access services such as the connection methods of the present application.

[00306] In some embodiments, M2M applications 20 and 20' may be used in conjunction with the disclosed systems and methods. The M2M applications 20 and 20' may include the applications that interact with the UE or gateway and may also be used in conjunction with other disclosed systems and methods.

[00307] In one embodiment, the logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48 may be hosted within a M2M service layer instance hosted by an M2M node, such as an M2M server, M2M gateway, or M2M device, as shown in Figure 49B. For example, the logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node

2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48 may comprise an individual service capability within the M2M service layer instance or as a sub-function within an existing service capability.

[00308] The M2M applications 20 and 20' may include applications in various industries such as, without limitation, transportation, health and wellness, connected home, energy management, asset tracking, and security and surveillance. As mentioned above, the M2M service layer, running across the devices, gateways, servers and other nodes of the system, supports functions such as, for example, data collection, device management, security, billing, location tracking/geofencing, device/service discovery, and legacy systems integration, and provides these functions as services to the M2M applications 20 and 20'.

[00309] Generally, the service layers 22 and 22' define a software middleware layer that supports value-added service capabilities through a set of Application Programming Interfaces (APIs) and underlying networking interfaces. Both the ETSI M2M and oneM2M architectures define a service layer. ETSI M2M's service layer is referred to as the Service Capability Layer (SCL). The SCL may be implemented in a variety of different nodes of the ETSI M2M architecture. For example, an instance of the service layer may be implemented within an M2M device (where it is referred to as a device SCL (DSCL)), a gateway (where it is referred to as a gateway SCL (GSCL)) and/or a network node (where it is referred to as a network SCL (NSCL)). The oneM2M service layer supports a set of Common Service Functions (CSFs) (i.e., service capabilities). An instantiation of a set of one or more particular types of CSFs is referred to as a Common Services Entity (CSE) which can be hosted on different types of network nodes (e.g. infrastructure node, middle node, application-specific node). The Third Generation Partnership Project (3GPP) has also defined an architecture for machine-type communications (MTC). In that architecture, the service layer, and the service capabilities it provides, are implemented as part of a Service Capability Server (SCS). Whether embodied in a DSCL, GSCL, or NSCL of the ETSI M2M architecture, in a Service Capability Server (SCS) of the 3GPP MTC architecture, in a CSF or CSE of the oneM2M architecture, or in some other node of a network, an instance of the service layer may be implemented as a logical entity (e.g., software, computer-executable instructions, and the like) executing either on one or more

standalone nodes in the network, including servers, computers, and other computing devices or nodes, or as part of one or more existing nodes. As an example, an instance of a service layer or component thereof may be implemented in the form of software running on a network node (e.g., server, computer, gateway, device or the like) having the general architecture illustrated in Figure 49C or Figure 49D described below.

[00310] Further, logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48 can be implemented as part of an M2M network that uses a Service Oriented Architecture (SOA) and/or a Resource-Oriented Architecture (ROA) to access services of the present application.

[00311] Figure 49C is a block diagram of an example hardware/software architecture of a M2M network node 30, such as an M2M device 18, an M2M gateway 14, an M2M server, or the like. The node 30 can execute or include logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48. The device 30 can be part of an M2M network as shown in Figure 49A-B or part of a non-M2M network. As shown in Figure 49C, the M2M node 30 may include a processor 32, non-removable memory 44, removable memory 46, a speaker/microphone 38, a keypad 40, a display, touchpad, and/or indicators 42, a power source 48, a global positioning system (GPS) chipset 50, and other peripherals 52. The node 30 may also include communication circuitry, such as a transceiver 34 and a transmit/receive element 36. It will be appreciated that the M2M node 30 may include any sub-combination of the foregoing elements while remaining consistent with an embodiment. This node may be a node that implements the functionality described herein.

[00312] The processor 32 may be a general purpose processor, a special purpose processor, a conventional processor, a digital signal processor (DSP), a plurality of microprocessors, one or more microprocessors in association with a DSP core, a controller, a microcontroller, Application Specific Integrated Circuits (ASICs), Field Programmable Gate Array (FPGAs) circuits, any other type of integrated circuit (IC), a state machine, and the like. In general, the processor 32 may execute computer-executable instructions stored in the memory (e.g., memory 44 and/or memory 46) of the node in order to perform the various required functions of the node. For example, the processor 32 may perform signal coding, data processing, power control, input/output processing, and/or any other functionality that enables the M2M node 30 to operate in a wireless or wired environment. The processor 32 may run application-layer programs (e.g., browsers) and/or radio access-layer (RAN) programs and/or other communications programs. The processor 32 may also perform security operations such as authentication, security key agreement, and/or cryptographic operations, such as at the access-layer and/or application layer for example.

[00313] As shown in Figure 49C, the processor 32 is coupled to its communication circuitry (e.g., transceiver 34 and transmit/receive element 36). The processor 32, through the execution of computer executable instructions, may control the communication circuitry in order to cause the node 30 to communicate with other nodes via the network to which it is connected. In particular, the processor 32 may control the communication circuitry in order to perform the transmitting and receiving steps described herein and in the claims. While Figure 49C depicts the processor 32 and the transceiver 34 as separate components, it will be appreciated that the processor 32 and the transceiver 34 may be integrated together in an electronic package or chip.

[00314] The transmit/receive element 36 may be configured to transmit signals to, or receive signals from, other M2M nodes, including M2M servers, gateways, device, and the like. For example, in an embodiment, the transmit/receive element 36 may be an antenna configured to transmit and/or receive RF signals. The transmit/receive element 36 may support various networks and air interfaces, such as WLAN, WPAN, cellular, and the like. In an embodiment, the transmit/receive element 36 may be an emitter/detector configured to transmit and/or receive IR, UV, or visible light signals, for example. In yet another embodiment, the transmit/receive element 36 may be configured to transmit and receive both RF and light signals. It will be appreciated that the transmit/receive element 36 may be configured to transmit and/or receive any combination of wireless or wired signals.

[00315] In addition, although the transmit/receive element 36 is depicted in Figure 49C as a single element, the M2M node 30 may include any number of transmit/receive elements 36. More specifically, the M2M node 30 may employ MIMO technology. Thus, in an embodiment, the M2M node 30 may include two or more transmit/receive elements 36 (e.g., multiple antennas) for transmitting and receiving wireless signals.

[00316] The transceiver 34 may be configured to modulate the signals that are to be transmitted by the transmit/receive element 36 and to demodulate the signals that are received by the transmit/receive element 36. As noted above, the M2M node 30 may have multi-mode capabilities. Thus, the transceiver 34 may include multiple transceivers for enabling the M2M node 30 to communicate via multiple RATs, such as UTRA and IEEE 802.11, for example.

[00317] The processor 32 may access information from, and store data in, any type of suitable memory, such as the non-removable memory 44 and/or the removable memory 46. For example, the processor 32 may store session context in its memory, as described above. The non-removable memory 44 may include random-access memory (RAM), read-only memory (ROM), a hard disk, or any other type of memory storage device. The removable memory 46 may include a subscriber identity module (SIM) card, a memory stick, a secure digital (SD) memory card, and the like. In other embodiments, the processor 32 may access information from, and store data in, memory that is not physically located on the M2M node 30, such as on a server or a home computer. The processor 32 may be configured to control visual indications on the display to reflect the status of the system or to obtain input from a user or display information to a user about capabilities or settings. A graphical user interface, which may be shown on the display, may be layered on top of an API to allow a user to interactively do functionality described herein.

[00318] The processor 32 may receive power from the power source 48, and may be configured to distribute and/or control the power to the other components in the M2M node 30. The power source 48 may be any suitable device for powering the M2M node 30. For example, the power source 48 may include one or more dry cell batteries (e.g., nickel-cadmium (NiCd), nickel-zinc (NiZn), nickel metal hydride (NiMH), lithium-ion (Li-ion), etc.), solar cells, fuel cells, and the like.

[00319] The processor 32 may also be coupled to the GPS chipset 50, which is configured to provide location information (e.g., longitude and latitude) regarding the current location of the M2M node 30. It will be appreciated that the M2M node 30 may acquire location

information by way of any suitable location-determination method while remaining consistent with an embodiment.

[00320] The processor 32 may further be coupled to other peripherals 52, which may include one or more software and/or hardware modules that provide additional features, functionality and/or wired or wireless connectivity. For example, the peripherals 52 may include various sensors such as an accelerometer, biometrics (e.g., fingerprint) sensors, an e-compass, a satellite transceiver, a digital camera (for photographs or video), a universal serial bus (USB) port or other interconnect interfaces, a vibration device, a television transceiver, a hands free headset, a Bluetooth® module, a frequency modulated (FM) radio unit, a digital music player, a media player, a video game player module, an Internet browser, and the like.

[00321] The node 30 may be embodied in other apparatuses or devices, such as a sensor, consumer electronics, a wearable device such as a smart watch or smart clothing, a medical or eHealth device, a robot, industrial equipment, a drone, a vehicle such as a car, truck, train, or airplane. The node 30 may connect to other components, modules, or systems of such apparatuses or devices via one or more interconnect interfaces, such as an interconnect interface that may comprise one of the peripherals 52. Alternately, the node 30 may comprise apparatuses or devices, such as a sensor, consumer electronics, a wearable device such as a smart watch or smart clothing, a medical or eHealth device, a robot, industrial equipment, a drone, a vehicle such as a car, truck, train, or airplane.

[00322] Figure 49D is a block diagram of an exemplary computing system 90 which may also be used to implement one or more nodes of an M2M network, such as an M2M server, gateway, device, or other node. Computing system 90 may comprise a computer or server and may be controlled primarily by computer readable instructions, which may be in the form of software, wherever, or by whatever means such software is stored or accessed. Computing system 90 can execute or include logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48. Computing system 90 can be an M2M device, user equipment, gateway, UE/GW or any other nodes including nodes of the mobile care network, service layer network application

provider, terminal device 18 or an M2M gateway device 14 for example. Such computer readable instructions may be executed within a processor, such as central processing unit (CPU) 91, to cause computing system 90 to do work. In many known workstations, servers, and personal computers, central processing unit 91 is implemented by a single-chip CPU called a microprocessor. In other machines, the central processing unit 91 may comprise multiple processors. Coprocessor 81 is an optional processor, distinct from main CPU 91, that performs additional functions or assists CPU 91. CPU 91 and/or coprocessor 81 may receive, generate, and process data related to the disclosed systems and methods for E2E M2M service layer sessions, such as receiving session credentials or authenticating based on session credentials.

[00323] In operation, CPU 91 fetches, decodes, and executes instructions, and transfers information to and from other resources via the computer's main data-transfer path, system bus 80. Such a system bus connects the components in computing system 90 and defines the medium for data exchange. System bus 80 typically includes data lines for sending data, address lines for sending addresses, and control lines for sending interrupts and for operating the system bus. An example of such a system bus 80 is the PCI (Peripheral Component Interconnect) bus.

[00324] Memories coupled to system bus 80 include random access memory (RAM) 82 and read only memory (ROM) 93. Such memories include circuitry that allows information to be stored and retrieved. ROMs 93 generally contain stored data that cannot easily be modified. Data stored in RAM 82 can be read or changed by CPU 91 or other hardware devices. Access to RAM 82 and/or ROM 93 may be controlled by memory controller 92. Memory controller 92 may provide an address translation function that translates virtual addresses into physical addresses as instructions are executed. Memory controller 92 may also provide a memory protection function that isolates processes within the system and isolates system processes from user processes. Thus, a program running in a first mode can access only memory mapped by its own process virtual address space; it cannot access memory within another process's virtual address space unless memory sharing between the processes has been set up.

[00325] In addition, computing system 90 may contain peripherals controller 83 responsible for communicating instructions from CPU 91 to peripherals, such as printer 94, keyboard 84, mouse 95, and disk drive 85.

[00326] Display 86, which is controlled by display controller 96, is used to display visual output generated by computing system 90. Such visual output may include text, graphics, animated graphics, and video. Display 86 may be implemented with a CRT-based video display, an LCD-based flat-panel display, gas plasma-based flat-panel display, or a touch-panel. Display

controller 96 includes electronic components required to generate a video signal that is sent to display 86.

[00327] Further, computing system 90 may contain communication circuitry, such as for example a network adaptor 97, that may be used to connect computing system 90 to an external communications network, such as network 12 of Figure 49A and Figure 49B, to enable the computing system 90 to communicate with other nodes of the network.

[00328] User equipment (UE) can be any device used by an end-user to communicate. It can be a hand-held telephone, a laptop computer equipped with a mobile broadband adapter, or any other device. For example, the UE can be implemented as the M2M terminal device 18 of Figures 49 A-B or the device 30 of Figure 49 C.

[00329] It is understood that any or all of the systems, methods, and processes described herein may be embodied in the form of computer executable instructions (i.e., program code) stored on a computer-readable storage medium which instructions, when executed by a machine, such as a node of an M2M network, including for example an M2M server, gateway, device or the like, perform and/or implement the systems, methods and processes described herein. Specifically, any of the steps, operations or functions described above, including the operations of the gateway, UE, UE/GW, or any of the nodes of the mobile core network, service layer or network application provider, may be implemented in the form of such computer executable instructions. Logical entities such as service layer 1602, data source 2218, DSAS 2202, DSAS-SF 2212, DSAS-SAE 2210, DSAS Manager 2206, DSAS API 2208, data storage node 2214, security manager 2302, manager/parser 2304, preprocessor 2306, storage filter 2308, analytics filter 2310, query store 2318, job table 2320, log store 2323, Stream ID Store 2321, SAE source 2314, query operators 2312, client 2216 and 3202, coordinator 3402, DSAS CSF 3702, DSAS hosting CSE 3704, data source AE 4402, DSAS Client AE 4502 and 4602, CSE 4504 and logical entities to produce interfaces such as the interfaces of Figures 47 and 48 may be embodied in the form of the computer executable instructions stored on a computer-readable storage medium. Computer readable storage media include both volatile and nonvolatile, removable and non-removable media implemented in any non-transitory (i.e., tangible or physical) method or technology for storage of information, but such computer readable storage media do not include signals. Computer readable storage media include, but are not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk

storage or other magnetic storage devices, or any other tangible or physical medium which can be used to store the desired information and which can be accessed by a computer.

[00330] In describing preferred embodiments of the subject matter of the present disclosure, as illustrated in the Figures, specific terminology is employed for the sake of clarity. The claimed subject matter, however, is not intended to be limited to the specific terminology so selected, and it is to be understood that each specific element includes all technical equivalents that operate in a similar manner to accomplish a similar purpose.

This written description uses examples to disclose the invention, including the best mode, and also to enable any person skilled in the art to practice the invention, including making and using any devices or systems and performing any incorporated methods. The patentable scope of the invention is defined by the claims, and may include other examples that occur to those skilled in the art. Such other examples are intended to be within the scope of the claims if they have elements that do not differ from the literal language of the claims, or if they include equivalent elements with insubstantial differences from the literal language of the claims.

What is Claimed:

1. An apparatus comprising a processor and a memory, the apparatus further including computer-executable instructions stored in the memory of the apparatus which, when executed by the processor of the apparatus, cause the apparatus to:

receive a data stream;
analyze, and filter the data stream; and
provide the filtered data stream to a node for further analysis

2. The apparatus of claim 1, wherein the apparatus is a Stream Forwarder (SF).

3. The apparatus of claim 1, wherein the node is a Stream Analytics Engine (SAE).

4. The apparatus of claim 3, wherein the Stream Analytics Engine is at a separate location from the apparatus.

5. The apparatus of claim 1, wherein the data stream is filtered according to one or more active queries.

6. The apparatus of claim 4, wherein the one or more active queries are obtained from a client.

7. The apparatus of claim 1, wherein the apparatus includes a service layer component.

8. The apparatus of claim 1, wherein the data stream comes from one or more Internet of Things (IoT) devices.

9. The apparatus of claim 1, wherein the data stream is parsed before filtering.

10. The apparatus of claim 9, wherein the parsed data stream is preprocessed before filtering.

11. An apparatus comprising a processor and a memory, the apparatus further including computer-executable instructions stored in the memory of the apparatus which, when executed by the processor of the apparatus, cause the apparatus to:

Receive a filtered data stream from a node; and
analyze the filtered data stream using at least one query.

12. The apparatus of claim 11, wherein the apparatus is a Stream Analytics Engine (SAE).

13. The apparatus of claim 11, wherein the node is a Stream Forwarder (SF).

14. The apparatus of claim 11, wherein the apparatus is at a different location from the node.

15. The apparatus of claim 14, wherein at least one query is received from at least one client.

16. The apparatus of claim 15, wherein the filtered data stream is analyzed using a Query Operator based on the query received by the client.

17. The apparatus of claim 14, wherein a response to the at least one query is sent to a client.

18. The apparatus of claim 11, wherein the apparatus includes a service layer component.

19. The apparatus of claim 11, wherein the apparatus includes at least one query operator that interacts with a Query Store.

20. The apparatus of claim 14, wherein an algorithm to process the at least one query is deployed by the client.

21. A method for use by an apparatus, wherein the apparatus comprises a processor and memory, and wherein the apparatuses further includes computer-executable instructions stored in the memory which, when executed by the processor, perform functions of a method comprising:

receiving a data stream;
analyzing, and filtering the data stream; and
providing the filtered data stream to a node for further analysis.

22. The method of claim 21, wherein the apparatus is Stream Forwarder

23. The method of claim 21, wherein the apparatus receives a data stream from IoT devices

24. The method of claim 21, wherein the node is a Stream Forwarder or Stream Analytics Engine

25. A method for use by an apparatus, wherein the apparatus comprises a processor and memory, and wherein the apparatuses further includes computer-executable instructions stored in the memory which, when executed by the processor, perform functions of a method comprising:

receiving a filtered data stream from a node; and
analyzing the filtered data stream using at least one query.

26. The method of claim 25, wherein the node is stream forwarder.

27. The method of claim 25, wherein the apparatus is stream analytics.

28. The method of claim 25, wherein the apparatus is connected to one or more additional apparatuses for distributed data stream analytics.




FIG. 1

4/11/2010	11:59:38	A	22.41618N	114.04299E	112.0	17.61
4/11/2010	11:59:39	A	22.41616N	114.04302E	113.0	14.09
4/11/2010	11:59:40	A	22.41611N	114.04301E	115.2	21.96
4/11/2010	11:59:41	A	22.41606N	114.04296E	115.6	32.28

FIG. 2






FIG. 3






FIG. 4C




FIG. 5

6/55




FIG. 6




FIG. 7

7/55




FIG. 8

8/55




FIG. 9

9/55




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14A




FIG. 14B




FIG. 15

15/55




FIG. 16




FIG. 17




FIG. 18




FIG. 19




FIG. 20

20/55





FIG. 21A




FIG. 21B




FIG. 22

23/55




FIG. 23




FIG. 24




FIG. 25




FIG. 26A




FIG. 26B




FIG. 27




FIG. 28

30/55




FIG. 29




FIG. 30A




FIG. 30B




FIG. 31




FIG. 32A




FIG. 32B




FIG. 33A




FIG. 33B





FIG. 33C

39/55






FIG. 35




FIG. 36




FIG. 37




FIG. 38




FIG. 39

FIG. 40

FIG. 41




FIG. 43




FIG. 42




FIG. 44




FIG. 45




FIG. 46A




FIG. 46B

4702

Query Output Panel

Select the Query from the dropdown:

Vehicle count / second at the intersection

Select the output format from the dropdown:

CSV file

Enter the time window for the query:

1 hour

Enter the maximum tolerable error (%):

0.01%

Submit

FIG. 47

Query Deployment Panel

Select the Query operation you are interested in performing from dropdown (add/delete/update/enable-disable):

Add new query

Select the query you want to update:

Vehicle count / second at intersection

Enter the Query Properties:

<Enter Query description>

<Enter Query Name>

Select stream(s) with their attributes required by the query from the dropdown table

<Enter Query Operator Name>

<Enter the namespace location where operator implementation will be uploaded>

<Upload the operator implementation>

Submit

Query Deployment Panel

Select the Query operation you are interested in performing from dropdown (add/delete/update/enable-disable):

Update existing query

Select the query you want to update:

Vehicle count / second at intersection

Select the query properties you want to update:

Query description

<Enter updated query description>

Submit

FIG. 48A

Query Deployment Panel

Select the Query operation you are interested in performing from dropdown (add/delete/update/enable-disable):

▼

Enable / Disable

Select the query you want to enable/disable:

▼

Vehicle count / second at intersection

The query is currently enabled. Do you want to disable it?

▼

Yes

Submit

Query Deployment Panel

Select the Query operation you are interested in performing from dropdown (add/delete/update/enable-disable):

▼

Delete

Select the query you want to delete:

▼

Vehicle count / second at intersection

Submit

FIG. 48B




FIG. 49A

53/55




FIG. 49B




FIG. 49C




FIG. 49D

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/029341

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04W4/00 G06F17/30 H04L12/28 H04L12/24 H04L29/06
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04W H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	WO 98/43178 A1 (INTEL CORP [US]) 1 October 1998 (1998-10-01) abstract page 2, line 7 - line 14 page 2, line 17 - line 24 page 3, line 9 - line 21 page 3, line 23 - page 4, line 3 page 4, line 21 - page 5, line 12 page 5, line 22 - page 6, line 4 page 6, line 27 - page 7, line 12 page 7, line 26 - page 9, line 7; claims 1,3-5,7,9,13,20; figures 1-3 ----- -/--	1-7, 9-22, 24-28 8,23



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

25 July 2017

Date of mailing of the international search report

31/07/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Todorut, Cosmin

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/029341

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 2 674 876 A1 (ALCATEL LUCENT [FR]) 18 December 2013 (2013-12-18) abstract paragraph [0001] - paragraph [0009] paragraphs [0013], [0015] - paragraph [0023] paragraph [0026] - paragraph [0034]; claims 1-3; figures 1-4 -----	1-28
X	US 2004/117157 A1 (BANTZ DAVID F [US] ET AL) 17 June 2004 (2004-06-17) abstract paragraph [0002] - paragraph [0004] paragraph [0012] - paragraph [0015] paragraph [0025] - paragraph [0034] paragraph [0038] - paragraph [0040]; claims 1-46; figures 1-6 -----	1-28
X	US 2015/271033 A1 (SRIVASTAVA ASHOK N [US] ET AL) 24 September 2015 (2015-09-24) abstract paragraph [0011] - paragraph [0018]; figures 1,2 paragraph [0022] - paragraph [0027]; figure 3 paragraph [0029] - paragraph [0047]; figure 4 paragraph [0050] - paragraph [0069]; figures 5a,5b,6 paragraph [0074] - paragraph [0077]; claims 1-20; figures 7a,7b -----	1-28
X	US 2014/330926 A1 (NANO OLIVIER [DE] ET AL) 6 November 2014 (2014-11-06) abstract paragraph [0001] - paragraph [0003] paragraphs [0006], [0017], [0021] - paragraph [0027]; figure 1 paragraphs [0031], [0035], [0036], [0039], [0041] paragraphs [0044], [0045] paragraphs [0047], [0048]; figure 2 paragraph [0051] - paragraph [0057]; figures 3,4 paragraph [0059] - paragraph [0062]; claims 1-20 ----- -/--	1-28

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2017/029341

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/280338 A1 (METZ CHRIS [US] ET AL) 18 September 2014 (2014-09-18) abstract paragraphs [0001], [0004] - paragraph [0006] paragraph [0014] - paragraph [0057]; figures 1-3 paragraph [0073] - paragraph [0103]; claims 1-27 -----	1-28
A	EP 2 801 935 A1 (ROCKWELL AUTOMATION TECH INC [US]) 12 November 2014 (2014-11-12) abstract paragraph [0006] - paragraph [0007] paragraph [0009] - paragraph [0012] paragraphs [0025], [0026], [0029], [0037], [0038] paragraph [0037] - paragraph [0045]; figure 4 paragraph [0048] - paragraph [0074]; claims 1-15; figures 6-14 -----	1-28
A	US 8 301 645 B1 (CROOK LAMONT ALLAN [US]) 30 October 2012 (2012-10-30) the whole document -----	1-28
A	AAZAM MOHAMMAD ET AL: "Fog Computing and Smart Gateway Based Communication for Cloud of Things", 2014 INTERNATIONAL CONFERENCE ON FUTURE INTERNET OF THINGS AND CLOUD, IEEE, 27 August 2014 (2014-08-27), pages 464-470, XP032704042, DOI: 10.1109/FICLOUD.2014.83 [retrieved on 2014-12-12] the whole document -----	1-28
A	Chris Campbell: "How to Process Google Data in Real Time with Azure Stream Analytics", BLUEGRANITE - Business Insights, 26 February 2015 (2015-02-26), page 1, XP055284658, Retrieved from the Internet: URL:https://www.blue-granite.com/blog/bid/403550/How-to-Process-Google-Data-in-Real-Time-with-Azure-Stream-Analytics [retrieved on 2016-06-29] the whole document -----	1-28

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2017/029341

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 9843178	A1	01-10-1998	AU 6571498 A 20-10-1998
		BR 9807896 A 21-03-2000	
		CN 1257593 A 21-06-2000	
		CN 1716940 A 04-01-2006	
		EP 1029286 A1 23-08-2000	
		US 6237031 B1 22-05-2001	
		WO 9843178 A1 01-10-1998	
EP 2674876	A1	18-12-2013	NONE
US 2004117157	A1	17-06-2004	NONE
US 2015271033	A1	24-09-2015	NONE
US 2014330926	A1	06-11-2014	NONE
US 2014280338	A1	18-09-2014	NONE
EP 2801935	A1	12-11-2014	CN 104142662 A 12-11-2014
		EP 2801935 A1 12-11-2014	
		US 2014337429 A1 13-11-2014	
		US 2016330291 A1 10-11-2016	
US 8301645	B1	30-10-2012	US 8301645 B1 30-10-2012
		US 2013060915 A1 07-03-2013	